

Two Embedding Theorems for Data with Equivalences under Finite Group Action

Fabian Lim*

Research Laboratory of Electronics, MIT, Cambridge, MA 02139, USA
flim@mit.edu

Abstract

There is recent interest in compressing data sets for non-sequential settings, where lack of obvious orderings on their data space, require notions of data equivalences to be considered. For example, Varshney & Goyal (DCC, 2006) considered multiset equivalences, while Choi & Szpankowski (IEEE Trans. IT, 2012) considered isomorphic equivalences in graphs. Here equivalences are considered under a relatively broad framework - finite-dimensional, non-sequential data spaces with equivalences under group action, for which analogues of two well-studied embedding theorems are derived: the Whitney embedding theorem and the Johnson-Lindenstrauss lemma. Only the canonical data points need to be carefully embedded, each such point representing a set of data points equivalent under group action. Two-step embeddings are considered. First, a group invariant is applied to account for equivalences, and then secondly, a linear embedding takes it down to low-dimensions. Our results require hypotheses on discriminability of the applied invariant, such notions related to *seperating invariants* (Dufresne, 2008), and *completeness* in pattern recognition (Kakarala, 1992).

Our first theorem shows that almost all such two-step embeddings can one-to-one embed the canonical part of a bounded, discriminable set of data points, if embedding dimension exceeds $2k$ whereby k is the box-counting dimension of the set closure of canonical data points. Our second theorem shows for k equal to the number of canonical points of a finite data set, a randomly sampled two-step embedding, preserves isometries (of the canonical part) up to factors $1 \pm \epsilon$ with probability at least $1 - \beta$, if the embedding dimension exceeds $(2 \log k + \log(1/\beta))/\alpha(\epsilon, \delta)$ for some function α , and δ is a positive constant capturing a certain discriminability property of the invariant. In both theorems, the value k is tied only to the canonical part, which may be significantly smaller than the ambient data dimension, up to a factor equal to the size the group.

*F. Lim recieved support from NSF Grant ECCS-1128226.




Figure 1: In (a), an exercise illustrating equivalences between three types of “non-conventional” data (for answers see below). In (b), accounting for data equivalences while performing embeddings.

1 Introduction

A discrete finite sequence is arguably the most *generic* mathematical representation for finite-dimensional data. However, of recent interest are data sets where it is unclear how to appropriately assign sequence orderings to the data space. For example, ranking data lives on a space of index subsets, which has no meaningful ordering [13]. Graphical data lives on a space of graph edges, and node labellings may be often irrelevant [9, 18]. *Quotient spaces* that describe matrix manifolds, *e.g.*, the *Grassman manifold*, have equivalence classes as elements [1].

We refer to such data sets as *non-sequential*, emphasizing the lack of ordering on their data space. For such sets, data compression becomes challenging. This is because we need to identify which seemingly different data points actually convey the same information. This is illustrated in Figure 1(a), whereby in each row, two (and only two) pictures are essentially the same (equivalent) but portrayed to appear different. Can you tell which two? The first row is designed to be an easy example, however the second row requires more time, and the third row is probably too difficult by human eye. These examples are not arbitrary, in fact they correspond to three previously studied “non-conventional” data models - the choice model [15], the *Ehrenfest* diffusion model (see [25], p. 5), and the graphical model (see [9, 18, 19]).

In this paper we extend low-dimensional *linear embedding* techniques [2, 3, 5, 6, 27], to the above mentioned non-sequential data models - more specifically, to finite-dimensional spaces where data equivalences result from a finite group action. We consider a two-step embedding process, illustrated in Figure 1(b). In the first step, we utilize a special function which produces the same output if two data sets are equivalent (under this group action); such a function, termed an *invariant*, accounts for data equivalence. Note however that the converse may not always hold, *i.e.*, two data sets producing the same output may not always be equivalent, such converses are related to *separating invariants* [14], and *completeness* in pattern recognition [17–19]. In the second step, a linear embedding is applied on the output of step one, to move the data to the low-dimensional space. The interest here is to obtain embedding guarantees, to support the use of such techniques as a kind of compression scheme. This has to be done with hypotheses on the discriminative power of the applied invariant, as an appropriate one-to-one embedding is not possible if the converse does not hold for any two data points of interest.

First row : two-three. Second row : one-four. Third row : one-two.

Main results: We extend two embedding theorems to finite-dimensional, non-sequential data spaces $\mathbb{R}[\mathcal{X}]$, discussed here for the case where the group \mathcal{G} acts by *permutation* action. Let \mathcal{R} denote a subset of $\mathbb{R}[\mathcal{X}]$, that contains canonical data points in $\mathbb{R}[\mathcal{X}]$, canonical under equivalence by action of \mathcal{G} . Then for a bounded set \mathcal{V} of data points (possibly *infinite*), assuming that the subset $\mathcal{V}_{\mathcal{R}}$ of canonical points (\mathcal{V} “projected” onto \mathcal{R}), are discriminable by the invariant (*i.e.*, satisfies the converse property), our extension (Theorem 3.1) of the Whitney’s embedding theorem shows that almost all such two-step embeddings can one-to-one embed $\mathcal{V}_{\mathcal{R}}$, if the embedding dimension exceeds $2k$ whereby k is the box-counting dimension of canonical points in set closure of $\mathcal{V}_{\mathcal{R}}$. For a *finite* set \mathcal{V} of data points, our extension (Theorem 3.2) of the Johnson-Lindenstrauss lemma shows that a randomly sampled linear embedding, preserves isometries up to factors $1 \pm \epsilon$ with probability at least $1 - \beta$, if the embedding dimension exceeds $(2 \log k + \log(1/\beta))/\alpha(\epsilon, \delta)$ for some function α , and δ is a positive constant that upper limits a to-be-defined undiscriminable fraction, between any two canonical points in $\mathcal{V}_{\mathcal{R}}$. The value k measuring the size of the set $\mathcal{V}_{\mathcal{R}}$ of canonical points, may be much smaller than that of the whole set \mathcal{V} , up to a factor $\#\mathcal{G}$ in group size. All proofs are simple and require little knowledge of invariant theory, facilitated by making obvious linear properties of invariants over a tensor space.

Significance of this work: These techniques are suited for database compression of non-sequential data, *e.g.*, DNA fragments, chemical molecular compositions, web-graph connections, record of intervallic events, etc. Here the models to admit any type of finite group (permutation) action - more general than specific cases considered in [9,26]. Extensions to any matrix group action seems feasible - to be pursued in future work. A synergistic relationship is developed between linear embeddings and (data) invariants, whereby this work can be viewed as an adaptation of invariants for low-dimensional data in high-dimensional ambient spaces. Provable guarantees are provided on the required *storage complexity* (embedding dimension), tied directly to the size of the data set. The invariant used in the second embedding step does not determine this complexity; it only needs to satisfy the discriminability hypothesis. While probabilistic data models are typically used in past related works [9,21,26], they are not required here. We discuss invariants with polynomial-time *computational complexity*, being at most mn^{ω} where m is embedding dimension, n is data-dimension of the model used, and $\omega \geq 1$. Compare with representation theoretic transform-type invariants (see [17–19]), where these methods require complexity of at least $\mathcal{O}((\#\mathcal{G})^2)$ to execute the fast transforms, a potentially large number if the group size $\#\mathcal{G}$ is huge ($\#\mathcal{G}$ may even be super-exponential in n for permutation groups, see [18], ch. 3 & 7).

More discussion on related prior work: Non-sequential data sets have been of interest for some years now, in pattern recognition [17], probability theory [13], machine learning [13,16,18], optimization [1,7], choice models [15], etc. Our interest in linear embeddings is due to the wealth of recent interest on this topic, *e.g.*, *compressed sensing* [6]. For invariant functions, the key area is invariant theory [12,14], though there exists other guises, *e.g.*, convex graphical invariants [7], triple-correlation [17,18], see also survey article [28]. One of their main applications of invariant theory is classification, and characterization of discriminative ability is of recent focus, see Dufresne’s Ph.D thesis [14]. For finite groups, a key result is that the set of all canonical points is in bijection with an *affine algebraic variety* corresponding to the *ideal of relations*, see [12], pp. 345-353; however the best known complexity bound is super-exponential in the number of data-dimensions n . For triple-correlation and equivalences under compact groups, Kakarala in his Ph.D thesis characterized the discriminative ability under certain conditions [17]. Kakarala uses representation theoretic techniques known as Tannaka-Krein duality. The difficulty in obtaining computationally efficient invariants with absolute discriminative ability, is appreciated by observing that even for the specific class of graphical invariants, a polynomial-time algorithm for *graph isomorphism* is still unknown for general graphs.

The work [26] is mainly an information theoretic study, for an efficient algorithm specialized for multisets see [21]. In [9] a very efficient $\mathcal{O}(\ell^2)$ algorithm specialized for compressing ℓ -node graphs is given, though their algorithm cannot be used as a graphical invariant. In both [9, 26], the dimension required for appropriate compression, is similar to that of our Johnson-Lindenstrauss lemma (Theorem 3.2) - there will be savings logarithmic $\log(\#\mathcal{G})$ in group size. For representation theoretic methods, partial labellings of graphical data is considered in [19].

For triple-correlations, Kakarala's proof in [17] is non-constructive, so an algorithm to invert an invariant function does not exist in general. However, invariant theory shows that the set of canonical points have a *manifold*, or *algebraic variety*, structure. Thus a possible future direction - inspired by compressed sensing - is to consider *manifold optimization* techniques (e.g., [1]) to perform inversion. In pattern recognition, correlation-type invariants are usually treated disparately from invariant theory, however they are related to polynomial functions from an invariant ring. However, do note that correlation invariants restrict to only *transitive* permutation group actions (where we say the data space is *homogeneous*). Also as Kondor pointed out [18], pp. 89-90, one needs to take care of Kakarala's notion of homogeneous spaces¹.

Organization: Section 2 touches on preliminaries, developing the type of invariants used in this work. Section 3 states the main results, on Whitney embedding (Subsection 3.2) and Johnson-Lindenstrauss (Subsection 3.3). Technical proofs are provided in Section 4.

Supplementary Material (SM-I & SM-II): For the sake of most readers who will not be familiar with both invariant theory, and representation theoretic analyses of correlation functions, two sets of supplementary materials are provided at the very end of this manuscript. Results from both these topics, alluded to throughout this text, are summarized in these materials.

2 Preliminaries

2.1 Finite-dimensional data \mathcal{G} -spaces: We assume some basic familiarity with *group theory*. Let \mathcal{G} denote a group, where h and g denote group elements. Let \mathcal{X} denote a set of a finite number n of elements, and x denotes an element of \mathcal{X} . Define a *permutation* action of group \mathcal{G} on the set \mathcal{X} , where $g(x)$ is the image of x under g , i.e., $g(x) \in \mathcal{X}$. This is a *left action*, i.e., for $h, g \in \mathcal{G}$ we have $(hg)(x) = h(g(x))$. A set \mathcal{X} endowed with such an action of \mathcal{G} is called a **\mathcal{G} -space**.

Let \mathbb{R} denote the set of real numbers. Let $\mathbb{R}[\mathcal{X}]$ denote a set of *real-valued* n -dimensional vectors, indexed over the set \mathcal{X} . Data points lie in this set. For $\mathbf{a} \in \mathbb{R}[\mathcal{X}]$, the element of \mathbf{a} indexed by x is written as a_x for all $x \in \mathcal{X}$. The space $\mathbb{R}[\mathcal{X}]$ (and therefore also the data) inherit the group action. If \mathbf{a}^g denotes the image of \mathbf{a} under g , i.e., $\mathbf{a}^g \in \mathbb{R}[\mathcal{X}]$, then we have $(\mathbf{a}^g)_{g(x)} = a_x$ for any $x \in \mathcal{X}$. By the left action of \mathcal{G} on \mathcal{X} given above, it follows that $\mathbf{a}^{hg} = (\mathbf{a}^g)^h$. While $\mathbb{R}[\mathcal{X}]$ can be identified with \mathbb{R}^n , the notation $\mathbb{R}[\mathcal{X}]$ emphasizes the group action. We illustrate using the following examples. Let e denote the group identity element of \mathcal{G} , and let $\#\mathcal{X}$ be the cardinality of \mathcal{X} .

Example. [Periodic data]: Let $\mathcal{X} = \{1, 2, \dots, n\}$. Let \mathcal{G} denote the n -th order cyclic group, i.e., $\mathcal{G} = \{e, g, g^2, \dots, g^{n-1}\}$, whereby \mathcal{G} acts on \mathcal{X} as follows: for the special element g , we have $g(i) = i + 1$ for $1 \leq i < n$, and $g(n) = 1$. This action is transitive.

Example. [Choice & graphical data]: Let \mathcal{X} be the set of size- ω subsets of $\{1, 2, \dots, \ell\}$, where the size $\#\mathcal{X} = \binom{\ell}{\omega}$. Let Sym_ℓ be the **symmetric group** (or the group of all permutations) on ℓ letters. Consider the group action of Sym_ℓ on \mathcal{X} , where for any $g \in \text{Sym}_\ell$, we have the image $g(\mathcal{V}) = \{g(i) : i \in \mathcal{V}\}$ for any $\mathcal{V} \in \mathcal{X}$. This action is transitive. The special case $\omega = 2$ corresponds to graphical data, as any graph is defined by the specification of $\binom{\ell}{2}$ edges.

¹ Kakarala's formulation of homogeneous spaces is different than that of Kondor (see Supplementary Material SM-II.1). Kondor points out that Kakarala's definition, in some cases, "do not model real-world problems as well". We tend to agree.

More generally, one would let \mathcal{G} act on \mathbb{R}^n as a *matrix group* - as in invariant theory [12, 14]. For simplicity, we focus only on permutation groups, which in fact covers all data models that apply for triple-correlation invariants [17–19].

2.2 \mathcal{G} -invariants with certain linearity properties: We provide bare minimal background on invariant theory. Those familiar with this material may find our presentation unconventional, as the material is discussed in the way that we feel best supports the exposition of our main results.

We build a tensor space using the vector space $\mathbb{R}[\mathcal{X}]$. For $\omega \geq 1$, let $\mathcal{X}^{\times\omega}$ denote the product set $\mathcal{X} \times \cdots \times \mathcal{X}$ between ω copies of \mathcal{X} . Then an ω -**array**, denoted $\llbracket b_{\mathbf{x}^{(1:\omega)}} \rrbracket$, has n^ω components $b_{\mathbf{x}^{(1:\omega)}}$ indexed over $\mathcal{X}^{\times\omega}$, *i.e.*, $\mathbf{x}^{(1:\omega)} \in \mathcal{X}^{\times\omega}$, where $\mathbf{x}^{(1:\omega)}$ denotes the ω -tuple $(x^{(1)}, \dots, x^{(\omega)})$. Let $\mathbb{R}[\mathcal{X}^{\times\omega}]$ denote the set of all ω -arrays over $\mathcal{X}^{\times\omega}$.

The tensor (outer) product between two elements \mathbf{a}, \mathbf{a}' in $\mathbb{R}[\mathcal{X}]$, denoted $\mathbf{a} \otimes \mathbf{a}'$, equals $(a_x \cdot a'_y)_{x,y \in \mathcal{X}}$. Multiple tensor products, denoted $\mathbf{a}^{(1)} \otimes \cdots \otimes \mathbf{a}^{(\omega)}$ for $\mathbf{a}^{(j)} \in \mathbb{R}[\mathcal{X}]$, $1 \leq j \leq \omega$, follow similarly. Now $\mathbf{a}^{(1)} \otimes \cdots \otimes \mathbf{a}^{(\omega)} \in \mathbb{R}[\mathcal{X}^{\times\omega}]$, by considering the ω -array $\llbracket a_{x^{(1)}} \cdots a_{x^{(\omega)}} \rrbracket$. In fact $\mathbb{R}[\mathcal{X}^{\times\omega}]$ is isomorphic to the space obtained by taking tensor products (between vector spaces) of ω copies of $\mathbb{R}[\mathcal{X}]$, see [11]. For this reason $\mathbb{R}[\mathcal{X}^{\times\omega}]$ is called a **tensor space**, where the dimension² of $\mathbb{R}[\mathcal{X}]$ equals n^ω . For any $\mathbf{a} \in \mathbb{R}[\mathcal{X}]$, we denote $\mathbf{a}^{\otimes\omega}$ to mean $\mathbf{a} \otimes \cdots \otimes \mathbf{a}$ with ω copies of \mathbf{a} .

We now explain how the tensor space $\mathbb{R}[\mathcal{X}^{\times\omega}]$ admits invariants. Firstly, $\mathcal{X}^{\times\omega}$ inherits the group action of \mathcal{G} on \mathcal{X} , where the image $g(\mathbf{x}^{(1:\omega)})$ of $\mathbf{x}^{(1:\omega)}$ under g equals $(g(x^{(1)}), \dots, g(x^{(\omega)}))$. This obtains an action of \mathcal{G} on $\mathbb{R}[\mathcal{X}^{\times\omega}]$, where for any $\llbracket b_{\mathbf{x}^{(1:\omega)}} \rrbracket \in \mathbb{R}[\mathcal{X}^{\times\omega}]$, the image $g(\llbracket b_{\mathbf{x}^{(1:\omega)}} \rrbracket)$ under g equals the ω -array $\llbracket b_{g^{-1}(\mathbf{x}^{(1:\omega)})} \rrbracket$ (meaning that its the $\mathbf{x}^{(1:\omega)}$ -th component of the image equals $b_{g^{-1}(\mathbf{x}^{(1:\omega)})}$). The previous action of \mathcal{G} on $\mathcal{X}^{\times\omega}$ induces an *equivalence relation* on $\mathcal{X}^{\times\omega}$, whereby $\mathbf{x}_1^{(1:\omega)}, \mathbf{x}_2^{(1:\omega)} \in \mathcal{X}^{\times\omega}$ are equivalent if there exists some g in \mathcal{G} that sends $g(\mathbf{x}_1^{(1:\omega)}) = \mathbf{x}_2^{(1:\omega)}$, see [25]. The equivalence classes here are called **\mathcal{G} -orbits** (on $\mathcal{X}^{\times\omega}$), denoted $\Omega_{\mathcal{G}}(\mathcal{X}^{\times\omega})$. Each \mathcal{G} -orbit $\Omega_{\mathcal{G}}(\mathcal{X}^{\times\omega})$ will be associated with a ω -array $\llbracket b_{\mathbf{x}^{(1:\omega)}} \rrbracket$, as follows

$$b_{\mathbf{x}^{(1:\omega)}} = \begin{cases} 1 & \text{if } \mathbf{x}^{(1:\omega)} \in \Omega_{\mathcal{G}}(\mathcal{X}^{\times\omega}), \\ 0 & \text{otherwise.} \end{cases} \quad (2.1)$$

Finally thinking of $\mathbb{R}[\mathcal{X}^{\times\omega}]$ as $\mathbb{R}^{(n^\omega)}$, define an **inner product** as

$$\langle \llbracket a_{\mathbf{x}^{(1:\omega)}} \rrbracket, \llbracket b_{\mathbf{x}^{(1:\omega)}} \rrbracket \rangle = \sum_{\mathbf{x}^{(1:\omega)} \in \mathcal{X}^{\times\omega}} a_{\mathbf{x}^{(1:\omega)}} \cdot b_{\mathbf{x}^{(1:\omega)}}, \quad (2.2)$$

and we can construct a **\mathcal{G} -invariant**, a function whose output is invariant under action of \mathcal{G} .

PROPOSITION 2.1. *Let \mathcal{G} be a finite group, with permutation action on data space \mathcal{X} . For some \mathcal{G} -orbit $\Omega_{\mathcal{G}}(\mathcal{X}^{\times\omega})$ on $\mathcal{X}^{\times\omega}$, where $\omega \geq 1$, let $f_{\Omega_{\mathcal{G}}(\mathcal{X}^{\times\omega})} : \mathbb{R}[\mathcal{X}^{\times\omega}] \rightarrow \mathbb{R}$ denote the mapping*

$$f_{\Omega_{\mathcal{G}}(\mathcal{X}^{\times\omega})} : \llbracket a_{\mathbf{x}^{(1:\omega)}} \rrbracket \mapsto \langle \llbracket a_{\mathbf{x}^{(1:\omega)}} \rrbracket, \llbracket b_{\mathbf{x}^{(1:\omega)}} \rrbracket \rangle \quad (2.3)$$

where $\llbracket b_{\mathbf{x}^{(1:\omega)}} \rrbracket$ is associated with $\Omega_{\mathcal{G}}(\mathcal{X}^{\times\omega})$ as in (2.1). Then $f_{\Omega_{\mathcal{G}}(\mathcal{X}^{\times\omega})}$ is a \mathcal{G} -invariant, *i.e.*, for any $\llbracket a_{\mathbf{x}^{(1:\omega)}} \rrbracket \in \mathbb{R}[\mathcal{X}^{\times\omega}]$ we have $f_{\Omega_{\mathcal{G}}(\mathcal{X}^{\times\omega})}(g(\llbracket a_{\mathbf{x}^{(1:\omega)}} \rrbracket)) = f_{\Omega_{\mathcal{G}}(\mathcal{X}^{\times\omega})}(\llbracket a_{\mathbf{x}^{(1:\omega)}} \rrbracket)$ for all $g \in \mathcal{G}$.

Proof. For brevity, write $\Omega = \Omega(\mathcal{X}^{\times\omega})$. Let $g \in \mathcal{G}$. By the earlier definition of the image of $\llbracket a_{\mathbf{x}^{(1:\omega)}} \rrbracket$ under g , the value $f_{\Omega}(g(\llbracket a_{\mathbf{x}^{(1:\omega)}} \rrbracket))$ is computed by summing the coefficients $a_{\mathbf{x}^{(1:\omega)}}$ supported over a subset \mathcal{V} , of the form $\mathcal{V} = \{g^{-1}(\mathbf{x}^{(1:\omega)}) : \mathbf{x}^{(1:\omega)} \in \Omega\}$. Since Ω is a \mathcal{G} -orbit, we may verify that \mathcal{V} is a $(g\mathcal{G}g^{-1})$ -orbit of $\mathcal{X}^{\times\omega}$, here $g\mathcal{G}g^{-1}$ is a group, $g\mathcal{G}g^{-1} = \{g\sigma g^{-1} : \sigma \in \mathcal{G}\}$. But $g\mathcal{G}g^{-1}$ is an automorphism of the group \mathcal{G} , hence $\mathcal{V} = \Omega$ and we conclude the result. \square

²If $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n$ is a basis of $\mathbb{R}[\mathcal{X}]$, then the n^ω tensors $\mathbf{e}_{\sigma(1)} \otimes \mathbf{e}_{\sigma(2)} \otimes \cdots \otimes \mathbf{e}_{\sigma(\omega)}$, for all $\sigma \in \text{Sym}_\omega$, consists a basis for the tensor space $\mathbb{R}[\mathcal{X}^{\times\omega}]$, see [11].

It is important to note that the \mathcal{G} -invariant (2.3) is *linear* in its domain $\mathcal{X}^{\times\omega}$. We extend these invariants to obtain the following linear \mathcal{G} -invariant $\mathcal{F}_\omega : \mathbb{R}[\mathcal{X}^{\times\omega}] \rightarrow \mathbb{R}^{\kappa_\omega}$ of main interest, by setting

$$\begin{aligned} \mathcal{F}_\omega : \llbracket a_{\mathbf{x}(1:\omega)} \rrbracket &\mapsto (z_1, z_2, \dots, z_{\kappa_\omega}), \\ z_i &= \#(\Omega_{\mathcal{G},i})^{-\frac{1}{2}} \cdot f_{\Omega_{\mathcal{G},i}}(\llbracket a_{\mathbf{x}(1:\omega)} \rrbracket), \end{aligned} \quad (2.4)$$

where κ_ω denotes the number of different \mathcal{G} -orbits on $\mathcal{X}^{\times\omega}$, numbered as $\Omega_{\mathcal{G},1}, \dots, \Omega_{\mathcal{G},\kappa_\omega}$, and $\omega \geq 1$. We propose to use (2.4) in the first embedding step (recall illustration Figure 1(b)).

ALGORITHM 2.1. \mathcal{G} -invariant (2.4) and embedding step one

- 1) for given data point $\mathbf{a} \in \mathbb{R}[\mathcal{X}]$, take the ω -th tensor power $\mathbf{a}^{\otimes\omega}$.
- 2) output the length- κ_ω vector $\mathcal{F}_\omega(\mathbf{a}^{\otimes\omega})$.

In the upcoming Section 3, the linearity of \mathcal{F}_ω will be exploited to connect with linear embedding theory. The normalization factor $\#(\Omega_{\mathcal{G},i})^{-\frac{1}{2}}$ w.r.t. orbit cardinality in (2.4) is so that \mathcal{F}_ω will have unity operator norm (to ensure stability).

But before going on to discussing embeddings, we clarify some properties of the invariants. Firstly, \mathcal{F}_ω has **polynomial complexity** of evaluation (in n for fixed ω), exactly n^ω . Next, the number of \mathcal{G} -orbits κ_ω over $\mathcal{X}^{\times\omega}$ determines the (dimension of the) range of \mathcal{F}_ω , and we call κ_ω the **invariant dimension**. We briefly discuss how to determine κ_ω . Let $\theta_{\mathcal{G},\mathcal{X}} : \mathcal{G} \rightarrow \mathbb{R}$, that satisfies

$$\theta_{\mathcal{G},\mathcal{X}}(g) = \#\{x \in \mathcal{X} : g(x) = x\} \quad (2.5)$$

for all $g \in \mathcal{G}$, i.e., the value $\theta_{\mathcal{G},\mathcal{X}}(g)$ equals the number of points in \mathcal{X} *fixed* by the permutation g in \mathcal{G} . The classical **Burnside lemma**, see e.g. [25], p. 106, allows us to determine κ_ω as follows

$$\kappa_\omega = \frac{1}{\#\mathcal{G}} \sum_{g \in \mathcal{G}} (\theta_{\mathcal{G},\mathcal{X}}(g))^\omega. \quad (2.6)$$

Note $\theta_{\mathcal{G},\mathcal{X}}(e) = \#\mathcal{X} = n$ for the identity element e .

Example. [Periodic data]: If \mathcal{G} equals the cyclic group on n letters, i.e., then $\theta_{\mathcal{G},\mathcal{X}}(g) = 0$ for all $g \neq e$. Since $\#\mathcal{G} = \#\mathcal{X} = n$, thus $\kappa_\omega = n^{\omega-1}$.

To simplify calculation of (2.5), one may use the fact that for any $g \in \mathcal{G}$, $\theta_{\mathcal{G},\mathcal{X}}(\sigma g \sigma^{-1}) = \theta_{\mathcal{G},\mathcal{X}}(g)$ for all $\sigma \in \mathcal{G}$, see the following example. There exists an equivalence relation on elements in \mathcal{G} , if we deem h equivalent with g if $h = \sigma g \sigma^{-1}$ for some $\sigma \in \mathcal{G}$, see [25], p. 81.

Example. [Graphical data]: For $\mathcal{G} = \text{Sym}_\ell$ with some integer ℓ , by the above relation there exists a bijection between equivalence classes, and the unordered partitions of ℓ , see [25], ch. 10. For example, we can express $\ell = 3$ as $1 + 1 + 1$, $2 + 1$, and 3 ; in the first partition three 1's appear, in the second partition one 1 appears and one 2 appears. One can use this bijection to show that $\theta_{\mathcal{G},\mathcal{X}}(g) = \{\# \text{ of 2's appearing}\} + \left(\frac{\# \text{ of 1's appearing}}{2}\right)$ for the partition corresponding to g .

REMARK 2.1. *In invariant theoretic terms, the \mathcal{G} -invariant \mathcal{F}_ω is equivalent to a generating set of the degree- ω homogeneous polynomials in the invariant ring, see supplementary material SM-I.1. Due to interest in applying invariants for classification, there is recent focus on studying minimal sets of invariants that discriminate between all data points, i.e., any $\mathbf{a}_1, \mathbf{a}_2 \in \mathbb{R}[\mathcal{X}]$ are never mistaken if $\mathbf{a}_1 \neq \mathbf{a}_2^g$ for all $g \in \mathcal{G}$, see [14] (Theorem SM-I.1). Unfortunately such powerful discriminability properties come at super-exponential complexity (Fact SM-I.1). Thus, it is meaningful to ask, for a given invariant \mathcal{F}_ω , what are the pairs of data points that it cannot discriminate. For \mathcal{F}_ω , this amounts to looking at an affine algebraic variety, see supplementary material SM-I.2. In particular for \mathcal{G} -spaces with transitive action, we can view*

\mathcal{F}_ω as a multi-correlation function (see SM-II.1), and relate to completeness results for the triple-correlation [17–19] (see SM-II.2).

3 Two Theorems on Low-Dimensional Linear Embeddings of Data-Invariants

3.1 Two-step linear embedding (Figure 1(b)): For some $\omega \geq 1$, first apply a \mathcal{G} -invariant in Algorithm 2.1 to place the data (some $\mathbf{a} \in \mathbb{R}[\mathcal{X}]$) in κ_ω dimensions. Next, use a linear map $\Phi : \mathbb{R}^{\kappa_\omega} \rightarrow \mathbb{R}^m$ to effect the dimension reduction, whereby $m < \min(\kappa_\omega, n)$. Specifically, compute

$$\Phi(\mathcal{F}_\omega(\mathbf{a}^{\otimes \omega})), \quad (3.7)$$

where for convenience $\Phi\mathcal{F}_\omega$ will stand for the concatenation of the map \mathcal{F}_ω followed by the map Φ . Clearly $\Phi\mathcal{F}_\omega$ is a \mathcal{G} -invariant, linear in the domain $\mathbb{R}[\mathcal{X}^{\times \omega}]$, and drops dimensions down to m .

We desire embeddings that map the data set, some $\mathcal{V} \subset \mathbb{R}[\mathcal{X}]$, onto the lower dimensional space in some injective manner. This is possibly only when the embedding dimension m is sufficiently large enough to accommodate the data set. The key here is that m can be much smaller than the ambient data dimension n , where m should really only be tied to the size of \mathcal{V} . Linear embeddings have been studied for when \mathcal{V} is a union of subspaces [5, 6, 20], and a smooth manifold [1, 4, 10]. Here we look at the case where \mathcal{V} comes from a finite-dimensional, non-sequential \mathcal{G} -spaces for finite groups \mathcal{G} . We derive analogues of two well-known embedding theorems, in this two-step setting that employs \mathcal{G} -invariants, for both the Whitney embedding theorem (Subsection 3.2) and the Johnson-Lindenstrauss lemma (Subsection 3.3).

3.2 How many dimensions are needed to embed non-sequential data? In Whitney embedding we consider \mathcal{V} to be a *bounded* subset of $\mathbb{R}[\mathcal{X}]$. The size of a bounded subset \mathcal{V} , will be measured by the *box-counting dimension*. For a bounded subset \mathcal{V} , we define: i) the *closure* $\bar{\mathcal{V}}$, and ii) the minimal number $N_\epsilon(\mathcal{V})$ of boxes with sides of length ϵ (in $\mathbb{R}[\mathcal{X}]$) required to cover \mathcal{V} , in a grid. The **box-counting dimension** is then defined as

$$\text{boxdim}(\mathcal{V}) = \lim_{\epsilon \rightarrow 0} \frac{\log N_\epsilon(\mathcal{V})}{-\log \epsilon} \quad (3.8)$$

if the limit exists. Roughly speaking, if $\text{boxdim}(\mathcal{V}) = d$, then $N_\epsilon(\mathcal{V}) \approx \epsilon^{-d}$. The **lower box-counting dimension**, denoted $\underline{\text{boxdim}}(\mathcal{V})$, is defined regardless by replacing the limit by \liminf .

From our two-step embedding (3.7), the map $\Phi\mathcal{F}_\omega$ cannot produce a one-to-one embedding for \mathcal{V} , since the linear tensor invariant \mathcal{F}_ω is not always one-to-one on ω -th tensor powers of \mathcal{V} . On the other hand, we do not care to discriminate between equivalent data points. Thus to state what is an appropriate or desirable embedding, we first define a canonical notion of elements in $\mathbb{R}[\mathcal{X}]$, of which we only discriminate between. To this end, define the following disjoint subsets of $\mathbb{R}[\mathcal{X}]$. For $\mathbf{a} \in \mathbb{R}[\mathcal{X}]$, we say \mathbf{a} is *un-fixable* if $\mathbf{a}^g \neq \mathbf{a}$ is satisfied for all $g \in \mathcal{G}$. Let \mathcal{R} denote an open set in $\mathbb{R}[\mathcal{X}]$. Let \mathcal{R} satisfy the following 3 properties: i) *all* elements of \mathcal{R} are un-fixable, ii) the $\#\mathcal{G}$ subsets $\{\mathbf{a}^g : \mathbf{a} \in \mathcal{R}\}$, one for each $g \in \mathcal{G}$, are *disjoint*, and iii) the union $\cup_{g \in \mathcal{G}} \{\mathbf{a}^g : \mathbf{a} \in \mathcal{R}\}$ contains *all* un-fixable elements in $\mathbb{R}[\mathcal{X}]$. There are exactly $\#\mathcal{G}$ disjoint³ open sets in $\mathbb{R}[\mathcal{X}]$ that satisfy the above properties. We call these open sets **fundamental regions**, and any one of them will give us our required canonical notion. For $\mathcal{V} \subset \mathbb{R}[\mathcal{X}]$, a set of canonical elements can be $\{\mathbf{a}^g \in \mathcal{R} : \mathbf{a} \in \mathcal{V}, g \in \mathcal{G}\}$, which we denote by $\mathcal{V}_\mathcal{R}$ for brevity. Our hypothesis on discriminability is now stated formally: a \mathcal{G} -invariant is said to be **discriminable** over a subset \mathcal{V} , if this function is one-to-one over $\mathcal{V}_\mathcal{R}$ where \mathcal{R} is any fundamental region (note that this definition does not depend on the choice of \mathcal{R}).

The following theorem is an analogue of Theorem 2.2. [24], for two-step linear embeddings (3.7)

³Since if \mathcal{R} satisfies these conditions, then $\{\mathbf{a}^g : \mathbf{a} \in \mathcal{R}\}$ for any $g \in \mathcal{G}$ also satisfies.

over finite dimensional \mathcal{G} -spaces.

THEOREM 3.1. *Let \mathcal{G} be a finite group. Let \mathcal{X} be a finite dimensional \mathcal{G} -space. For some $\omega \geq 1$, let \mathcal{F}_ω be the \mathcal{G} -invariant in (2.4). Let \mathcal{R} be any fundamental region.*

Let \mathcal{V} denote the data set, $\mathcal{V} \subset \mathbb{R}[\mathcal{X}]$, and assume \mathcal{V} is bounded. Assume \mathcal{F}_ω is discriminable over \mathcal{V} , and let $k = \text{boxdim}(\overline{\mathcal{V}_\mathcal{R}})$, where we assume this limit k exists.

Let Φ be a linear map, that drops dimension from κ_ω to m . Then if $m > 2k$, then almost all such linear maps, the concatenated map $\Phi\mathcal{F}_\omega$ will be discriminable over \mathcal{V} .

The two-step linear embedding (3.7) with embedding dimension twice that of the data set, is guaranteed to appropriately embed a data set \mathcal{V} as long as the linear tensor \mathcal{G} -invariant \mathcal{F}_ω is discriminable over \mathcal{V} .

We make three comments on Theorem 3.1, starting with storage complexity. In its original version [24] for sequence data spaces, the value k , is taken as the box-counting dimension of the (closure of the) *whole* data set \mathcal{V} . Here however, $k = \text{boxdim}(\overline{\mathcal{V}_\mathcal{R}})$, which may be as much as $\#\mathcal{G}$ times lower than $\text{boxdim}(\overline{\mathcal{V}})$. This “factor of $\#\mathcal{G}$ savings” is intuitively expected, as we only need to differentiate between canonical elements in \mathcal{R} . Similar savings have been reported in other works [9, 26].

Secondly the computational complexity of evaluating $\Phi\mathcal{F}_\omega$ is exactly mn^ω , polynomial in data dimension n (for fixed m, ω). Each coordinate of $\Phi\mathcal{F}_\omega$ is obtained by a weighted average of linear functions $f_{\Omega_{\mathcal{G},i}(\mathcal{X}^{\times\omega})}$, $1 \leq i \leq \kappa_\omega$.

Thirdly the linearity of $\Phi\mathcal{F}_\omega$ may be exploited to reduce computation. For example in [19], Kondor et. al. used a subspace of $\mathbb{R}[\text{Sym}_\ell]$ to represent⁴ graphical data on ℓ nodes, a (Sym_ℓ) -space where $n = \ell!$, see [18]. Now if the data lives in a k -dimensional subspace \mathcal{V} of $\mathbb{R}[\mathcal{X}]$, $k < n$, let $A : \mathbb{R}^k \rightarrow \mathcal{V}$ be a linear map onto \mathcal{V} . Then the tensor product map $A^{\otimes\omega} : \mathbb{R}^{k^\omega} \rightarrow \mathcal{V}^{\otimes\omega}$, where $\mathcal{V}^{\otimes\omega} \subset \mathbb{R}[\mathcal{X}^{\times\omega}]$, is *linear* in its domain \mathbb{R}^{k^ω} . Now the concatenated map from \mathbb{R}^{k^ω} to \mathbb{R}^m will be $\Phi\mathcal{F}_\omega A^{\otimes\omega}$, where each coordinate is obtained by a map obtained from a weighted average of functions $f_{\Omega_{\mathcal{G},i}(\mathcal{X}^{\times\omega})} A^{\otimes\omega}$, $1 \leq j \leq \kappa_\omega$, and this map is linear (and can be evaluated in κ^ω operations. Hence, the total evaluation complexity of \mathbb{R}^{k^ω} to \mathbb{R}^m equals mk^ω , where again k is the data dimension. In the above example where $\mathcal{X} = \text{Sym}_\ell$, we have $k = \binom{\ell}{2}$, so the complexity equals $\mathcal{O}(m\ell^{2\omega})$, which (for fixed m, ω) is polynomial in the number of nodes ℓ .

3.3 How many dimensions are needed to preserve isometries of non-sequential data?

Theorem 3.1 does not provide any notion of distance isometries under embedding, important for certain “sketching”-type applications. An important result for isometry preservation is the Johnson-Lindenstrauss lemma. In this part, the data set \mathcal{V} will be assumed to contain a *finite* number of discrete points in $\mathbb{R}[\mathcal{X}]$. Also here, we state the discriminability hypothesis slightly differently. By 2-norm $\|\cdot\|_2$ on elements in $\mathbb{R}[\mathcal{X}^{\times\omega}]$, we mean the norm

$$\| [b_{\mathbf{x}(1:\omega)}] \|_2 = \sqrt{\sum_{\mathbf{x}(1:\omega) \in \mathcal{X}^{\times\omega}} b_{\mathbf{x}(1:\omega)}^2}. \quad (3.9)$$

as if we were treating $\mathbb{R}[\mathcal{X}^{\times\omega}]$ as $\mathbb{R}^{(n^\omega)}$. Assuming that \mathcal{F}_ω is discriminable over \mathcal{V} , there must exist some constant $\delta < 1$, such that if for any $\mathbf{a}_1, \mathbf{a}_2 \in \mathcal{V}_\mathcal{R}$, where \mathcal{R} is any fundamental region, we have

$$\| A_{\mathcal{F}_\omega}(\mathbf{a}_1^{\otimes\omega} - \mathbf{a}_2^{\otimes\omega}) \|_2^2 \leq \delta \cdot \| \mathbf{a}_1^{\otimes\omega} - \mathbf{a}_2^{\otimes\omega} \|_2^2, \quad (3.10)$$

where $A_{\mathcal{F}_\omega} : \mathbb{R}[\mathcal{X}^{\times\omega}] \rightarrow \mathbb{R}[\mathcal{X}^{\times\omega}]$ is the orthogonal projection onto the kernel of \mathcal{F}_ω . That is for canonical elements $\mathbf{a}_1, \mathbf{a}_2 \in \mathcal{V}_\mathcal{R}$, the constant δ captures the maximal fraction of “energy” of the

⁴Kondor et. al. represented each data corresponding to edge $\{i, j\}$, in a redundant fashion using multiple coefficients a_x of $\mathbf{a} \in \mathbb{R}[\text{Sym}_\ell]$, for all x that send $\{\ell - 1, \ell\}$ to $\{i, j\}$.

error $\mathbf{a}_1^{\otimes \omega} - \mathbf{a}_2^{\otimes \omega}$ in the kernel of \mathcal{F}_ω .

The following theorem is an analogue of (the most basic form of) the Johnson-Lindenstrauss lemma, for two-step linear embeddings (3.7) over finite \mathcal{G} -spaces. The result is stated for the case where the coefficients of Φ are sampled from the normal distribution. However as in many works [2, 3, 6, 22], extensions to more general distributions should not be too difficult.

THEOREM 3.2. *We take $\mathcal{X}, \mathcal{G}, \mathbb{R}[\mathcal{X}]$ and \mathcal{R} as defined in Theorem 3.1. Let \mathcal{V} contain a finite number of discrete points in $\mathbb{R}[\mathcal{X}]$. Let $k = \#\mathcal{V}_\mathcal{R}$. For some $\omega \geq 1$, assume \mathcal{F}_ω is discriminable over \mathcal{V} , and that the constant $\delta < 1$ satisfies (3.10). Assume that the size $m \times \kappa_\omega$ linear map Φ , has coefficients independently sampled from a normal distribution with variance $1/m$. Then with probability at least $1 - \beta$, if the embedding dimension m of the map Φ exceeds*

$$\frac{2 \log k + \log(1/\beta)}{\alpha((\epsilon - \delta)/(1 - \delta))} \quad (3.11)$$

where $\alpha(y) = y^2 - y^3$ for any $y \in \mathbb{R}$, we will have for any $\mathbf{a}_1, \mathbf{a}_2 \in \mathcal{V}$, $\mathbf{a}_1 \neq \mathbf{a}_2$, the following isometries

$$\|\Phi \mathcal{F}_\omega(\mathbf{a}_1^{\otimes \omega} - \mathbf{a}_2^{\otimes \omega})\|_2^2 \begin{cases} \leq (1 + \epsilon) \cdot \|\mathbf{b}_1^{\otimes \omega} - \mathbf{b}_2^{\otimes \omega}\|_2^2, \\ \geq (1 - \epsilon) \cdot \|\mathbf{b}_1^{\otimes \omega} - \mathbf{b}_2^{\otimes \omega}\|_2^2, \end{cases} \quad (3.12)$$

for any positive $\epsilon > \delta$, and canonical elements $\mathbf{b}_1, \mathbf{b}_2$ (where $\mathbf{b}_1 = \mathbf{a}_1^{g_1}$ and $\mathbf{b}_2 = \mathbf{a}_2^{g_2}$ for some $g_1, g_2 \in \mathcal{G}$ such that $\mathbf{b}_1, \mathbf{b}_2 \in \mathcal{V}_\mathcal{R}$).

The factor ϵ in (3.12) should not be too close to the constant δ in (3.10) - this increases the required value for m (it affects the denominator of (3.11)). Also again thanks to \mathcal{G} -invariance, the (potential) “factor of $\#\mathcal{G}$ ” savings appear in k (here $k = \#\mathcal{V}_\mathcal{R}$ not $k = \#\mathcal{V}$). Do note there is a difference how these savings impact the embedding dimension m ; unlike the previous Theorem 3.1 where the factor of $\#\mathcal{G}$ impacts m multiplicatively (seen from the required assumption $m > 2k$), in Theorem 3.2 this factor impacts m logarithmically (seen from (3.11)). Also as seen from (3.12), the isometries are measured in the tensor space $\mathbb{R}[\mathcal{X}^{\times \omega}]$ (not in the data space $\mathbb{R}[\mathcal{X}]$). If one desires isometries in the original space, one requires some equivalence between the 2-norms of both spaces $\mathbb{R}[\mathcal{X}]$ and $\mathbb{R}[\mathcal{X}^{\times \omega}]$, not addressed here.

The next section provides technical proofs for the Theorems 3.1 and 3.2.

4 Technical Proofs

4.1 Proof of Theorem 3.1: The proof follows relatively closely with [24], though the consideration of \mathcal{G} -invariants allow certain simplifications, also see [5].

First some new notation. For any $\mathbf{a} \in \mathbb{R}^n$, for some positive integer n , we denote $\mathcal{B}_n(\mathbf{a}, \epsilon)$ to be the n -dimensional ball of radius ϵ , centered at \mathbf{a} . For any map, sometimes denoted A here, for any set \mathcal{V} that lies in the range of A , we shall use $A^{-1}(\mathcal{V})$ to denote the *pre-image* of \mathcal{V} . For any $\mathcal{V} \subset \mathbb{R}^n$ for any n , we denote the volume of \mathcal{V} as $\text{vol}(\mathcal{V})$. We will need the following two lemmas, simplified from [24]. For convenience, the lemma proofs are reproduced in Appendix A

LEMMA 4.1. (c.f. LEMMA 4.2, [24]) *For some positive integers r, m , $m \leq r$, let A be some surjective linear map from \mathbb{R}^r to \mathbb{R}^m . Let $\sigma > 0$ be a smallest singular value of A , obtained from any matrix form for A . Then for any $\epsilon > 0$*

$$\frac{\text{vol}(A^{-1}(\mathcal{B}_m(\epsilon)) \cap \mathcal{B}_r(\delta))}{\text{vol}(\mathcal{B}_r(\delta))} < 2^{r/2} \cdot \left(\frac{\epsilon}{\sigma \delta}\right)^m, \quad (4.13)$$

where $\mathcal{B}_r(\epsilon)$ and $\mathcal{B}_m(\epsilon)$ are respectively r - and m -dimensional balls centered at the origin.

LEMMA 4.2. (c.f. LEMMA 4.3, [24]) *Let \mathcal{V} be a bounded subset of \mathbb{R}^n , with $k = \text{boxdim}(\overline{\mathcal{V}})$, and*

we assume this limit k exists. Let ρ_1, \dots, ρ_r be r number of Lipschitz maps from \mathbb{R}^n to \mathbb{R}^m . Further assume that for each $\mathbf{a} \in \mathcal{V}$, the linear map $A : \mathbb{R}^r \rightarrow \mathbb{R}^m$ described by the matrix $[\rho_1(\mathbf{a}), \dots, \rho_r(\mathbf{a})]$, is surjective.

For each $\boldsymbol{\beta} \in \mathbb{R}^r$ with bounded 2-norm, $\boldsymbol{\beta} = [\beta_1, \dots, \beta_r]$, define $\rho_{\boldsymbol{\beta}} = \sum_{i=1}^r \beta_i \rho_i$. Then for almost every such bounded $\boldsymbol{\beta}$, the preimage $\rho_{\boldsymbol{\beta}}^{-1}(\mathbf{0})$ of the map $\rho_{\boldsymbol{\beta}}$ w.r.t. the single point $\mathbf{0}$, has lower box-counting dimension at most $k - m$. If $k > m$, then $\rho_{\boldsymbol{\beta}}^{-1}(\mathbf{0})$ is empty for almost every $\boldsymbol{\beta}$.

Proof. [Proof of Theorem 3.1]

We begin by making a connection with Lemma 4.2, first specifying for some positive integers n_2, r , the Lipschitz maps ρ_1, \dots, ρ_r (where each $\rho_i : \mathbb{R}^{n_2} \rightarrow \mathbb{R}^m$), and vectors $\boldsymbol{\beta}$ in \mathbb{R}^r . Note, here n_2 replaces n in Lemma 4.2.

The domain \mathbb{R}^{n_2} , where $n_2 = n^w$, is identified with $\mathbb{R}[\mathcal{X}^{\times \omega}]$, and we set the maps $\rho_i : \mathbb{R}[\mathcal{X}^{\times \omega}] \rightarrow \mathbb{R}^m$ as

$$\rho_{i+m(j-1)} : \llbracket a_{\mathbf{x}(1:\omega)} \rrbracket \mapsto \#(\Omega_{\mathcal{G},j})^{-\frac{1}{2}} \cdot f_{\Omega_{\mathcal{G},j}}(\llbracket a_{\mathbf{x}(1:\omega)} \rrbracket) \cdot \mathbf{e}_i \quad (4.14)$$

using the 1-Lipschitz functions $f_{\Omega_{\mathcal{G},j}}$ appearing in (2.4), for all $1 \leq i \leq m$, $1 \leq j \leq \kappa_{\omega}$, and where $\mathbf{e}_1, \dots, \mathbf{e}_m$ constitute any basis of \mathbb{R}^m . Thus here $r = m\kappa_{\omega}$, and we associate each vector $\boldsymbol{\beta}$ in $\mathbb{R}^{m\kappa_{\omega}}$ with the linear map $\Phi : \mathbb{R}^{\kappa_{\omega}} \rightarrow \mathbb{R}^m$, where $\boldsymbol{\beta}$ is formed by column-wise stacking of the coefficients from the matrix representation of Φ . Under these associations, it becomes clear that the map $\rho_{\boldsymbol{\beta}} : \mathbb{R}[\mathcal{X}^{\times \omega}] \rightarrow \mathbb{R}^m$ in the statement of Lemma 4.2, equals $\Phi \mathcal{F}_{\omega}$.

Let $\mathcal{V}^{(2)} = \{\mathbf{a}_1^{\otimes \omega} - \mathbf{a}_2^{\otimes \omega} : \mathbf{a}_1, \mathbf{a}_2 \in \overline{\mathcal{V}_{\mathcal{R}}}, \mathbf{a}_1 \neq \mathbf{a}_2\}$, i.e., $\mathcal{V}^{(2)}$ is (homomorphic) to the set of non-equal pairs of $\overline{\mathcal{V}_{\mathcal{R}}}$. We want to apply Lemma 4.2 with $\mathcal{V}^{(2)}$ replacing \mathcal{V} , with $2k$ replacing k (since $\text{boxdim}(\mathcal{V}^{(2)}) \leq 2k$). If the lemma applies, this shows one-to-one mapping on $\mathcal{V}_{\mathcal{R}}$, which proves the theorem. To do so, we need to show that for each $\llbracket a_{\mathbf{x}(1:\omega)} \rrbracket \in \mathcal{V}_2$, the linear map $A : \mathbb{R}^{m\kappa_{\omega}} \rightarrow \mathbb{R}^m$ as described in the statement of Lemma 4.2, is surjective. This will follow from the hypothesis that \mathcal{F}_{ω} is discriminable over \mathcal{V} , which implies that for each $\llbracket a_{\mathbf{x}(1:\omega)} \rrbracket \in \mathcal{V}_2$, there exists some function $f_{\Omega_{\mathcal{G},j}}$, $1 \leq j \leq \kappa_{\omega}$, such that $f_{\Omega_{\mathcal{G},j}}(\llbracket a_{\mathbf{x}(1:\omega)} \rrbracket) \neq 0$. By the association of A with the matrix $[\rho_1(\llbracket a_{\mathbf{x}(1:\omega)} \rrbracket), \dots, \rho_{m\kappa_{\omega}}(\llbracket a_{\mathbf{x}(1:\omega)} \rrbracket)]$, from (4.14) we conclude that since $f_{\Omega_{\mathcal{G},j}}(\llbracket a_{\mathbf{x}(1:\omega)} \rrbracket) \neq 0$ for some j , the map A will indeed be surjective. Thus the result is proved. \square

The key to the proof is the discriminability hypothesis. The important point is that does not impact embedding dimension m ; here m is tied directly to data size (tied to $k = \text{boxdim}(\overline{\mathcal{V}_{\mathcal{R}}})$). We also point out that while Sauer et. al. discuss more generalized versions of Lemmas 4.1 and 4.2 that do not require surjectivity of A (see [24], Lemma 4.6), these generalizations are not useful here. This is because as our proof of Theorem 3.1 reveals, the map A is either surjective (in the case discriminability holds) or otherwise the zero-map (in the case discriminability does not hold).

4.2 Proof of Theorem 3.2: The proof here also follows with simple modifications, by appropriately incorporating discriminability notions. Standard concentration results, such as the following one, will be useful (for convenience, its proof is reproduced in Appendix A).

LEMMA 4.3. (c.f., [2, 3]) Let \mathbf{A} be an $m \times \ell$ random matrix, whose matrix entries are standard normal RVs. Let the rows of \mathbf{A} be independent. Then for any $\mathbf{x} \in \mathbb{R}^{\ell}$, for any $\epsilon > 0$ we have

$$\Pr \left\{ \left| \|(1/\sqrt{m}) \cdot \mathbf{A}\mathbf{x}\|_2^2 - \|\mathbf{x}\|_2^2 \right| \leq \epsilon \right\} \geq 1 - 2e^{-\frac{m}{4}(\epsilon^2 - \epsilon^3)} \quad (4.15)$$

The proof of Theorem 3.2 given below will follow for other (row independent) distributions of \mathbf{A} , if probabilistic inequalities similar to (4.15) are available. Indeed they are for many other of distributions, see e.g., [2, 3, 27]. We do not go further into detail since this component is not our main focus. We use Lemma 4.3 to prove our second main theorem.

Proof. [Proof of Theorem 3.2]

It suffices to show the result for pairs $\mathbf{a}_1, \mathbf{a}_2 \in \mathcal{V}_{\mathcal{R}}$, $\mathbf{a}_1 \neq \mathbf{a}_2$, of canonical elements, since the LHS of (3.12) remains constant when replacing $\mathbf{a}_1, \mathbf{a}_2$ with $\mathbf{b}_1, \mathbf{b}_2$. For Φ uniformly sampled (recall lemma statement) as $\mathbf{A} = \Phi$, the probability that

$$\|\Phi \mathcal{F}_{\omega}(\mathbf{a}_1^{\otimes \omega} - \mathbf{a}_2^{\otimes \omega})\|_2^2 \begin{cases} \leq (1 + \epsilon) \cdot \|\mathcal{F}_{\omega}(\mathbf{a}_1^{\otimes \omega} - \mathbf{a}_2^{\otimes \omega})\|_2^2, \\ \geq (1 - \epsilon) \cdot \|\mathcal{F}_{\omega}(\mathbf{a}_1^{\otimes \omega} - \mathbf{a}_2^{\otimes \omega})\|_2^2, \end{cases} \quad (4.16)$$

holds for all $\binom{k}{2} < k^2/2$ pairs whereby $\mathbf{a}_1, \mathbf{a}_2 \in \mathcal{V}_{\mathcal{R}}$, is at least $1 - k^2 \cdot e^{-\frac{m}{4}(\epsilon^2 - \epsilon^3)}$. Here we used Lemma 4.3 for each $\mathbf{x} = \Phi \mathcal{F}_{\omega}(\mathbf{a}_1^{\otimes \omega} - \mathbf{a}_2^{\otimes \omega})$, $\mathbf{x} \in \mathbb{R}^m$. Comparing (4.16) with (3.12), the norm $\|\cdot\|_2$ on the RHS needs to be applied on the $\mathbb{R}[X^{\times \omega}]$, not \mathbb{R}^m . Recall from its definition, see (2.4), that \mathcal{F}_{ω} is 1-Lipschitz and linear in $\mathbb{R}[\mathcal{X}^{\times \omega}]$, so the upper bound follows as

$$\|\mathcal{F}_{\omega}(\mathbf{a}_1^{\otimes \omega} - \mathbf{a}_2^{\otimes \omega})\|_2^2 \leq \|(\mathbf{a}_1^{\otimes \omega} - \mathbf{a}_2^{\otimes \omega})\|_2^2.$$

For the lower bound, we use the hypothesis \mathcal{F}_{ω} is δ -discriminable over \mathcal{V} , where for the orthogonal projection $A_{\mathcal{F}_{\omega}} : \mathbb{R}[\mathcal{X}^{\times \omega}] \rightarrow \mathbb{R}[\mathcal{X}^{\times \omega}]$ onto the kernel of \mathcal{F}_{ω} , see (3.10), we have

$$\begin{aligned} \|\mathcal{F}_{\omega}(\mathbf{a}_1^{\otimes \omega} - \mathbf{a}_2^{\otimes \omega})\|_2^2 + \delta \cdot \|\mathbf{a}_1^{\otimes \omega} - \mathbf{a}_2^{\otimes \omega}\|_2^2 &\geq \|\mathcal{F}_{\omega}(\mathbf{a}_1^{\otimes \omega} - \mathbf{a}_2^{\otimes \omega})\|_2^2 + \|A_{\mathcal{F}_{\omega}}(\mathbf{a}_1^{\otimes \omega} - \mathbf{a}_2^{\otimes \omega})\|_2^2 \\ &= \|(\mathbf{a}_1^{\otimes \omega} - \mathbf{a}_2^{\otimes \omega})\|_2^2, \end{aligned} \quad (4.17)$$

equality following because both \mathcal{F}_{ω} and $A_{\mathcal{F}_{\omega}}$ project onto “orthogonal”⁵ spaces, which implies

$$\|\mathcal{F}_{\omega}(\mathbf{a}_1^{\otimes \omega} - \mathbf{a}_2^{\otimes \omega})\|_2^2 \geq (1 - \delta) \cdot \|\mathbf{a}_1^{\otimes \omega} - \mathbf{a}_2^{\otimes \omega}\|_2^2.$$

Using this in (4.16) and rearranging $(1 - \epsilon)(1 - \delta)$, this proves that (3.12) is satisfied with required probability, for constant $\epsilon(1 - \delta) + \delta > \epsilon$ (the strict inequality follows since $\delta > 0$). The statement of the proposition will satisfy for some probability $\beta > k^2 \cdot e^{-\frac{m}{4}(\epsilon^2 - \epsilon^3)}$, and rescaling the ϵ term used here. \square

The linearity of the \mathcal{G} -invariant \mathcal{F}_{ω} is very useful for deriving the lower bound (4.17), which admitted the use of orthonormality concepts. It is also useful for deriving the upper bound, since it made it easy to check that \mathcal{F}_{ω} is 1-Lipschitz. We are now done with the proofs of both main results.

REMARK 4.1. *For finite groups, there always exists an invariant satisfying the discriminability hypothesis [14] (albeit with super-exponential complexity, see Theorem SM-I.1 and Fact SM-I.1). However from an embedding complexity standpoint, for any non-sequential data set, (theoretically) one can always find a two-step embedding meeting the guarantees in both Theorems 3.1 and 3.2.*

Also, the canonical points in any fundamental region \mathcal{R} , have a manifold structure within an algebraic variety (see supplementary material SM-I.2). Hence an interesting future direction is to connect with manifold learning techniques (e.g., [1]).

5 Conclusion

We present a new extension of linear embeddings for non-sequential data, providing two theorems in the vein of Whitney embedding and the Johnson-Lindenstrauss lemma. We show that accounting for data equivalences can provide savings in embedding dimension up to a factor equal to the size of the invariance group (the savings is logarithmic in the second theorem). The extension was fairly simple, and we appeal to certain linearity properties of invariants.

Acknowledgment

The author thanks J. Z. Sun for discussions and his reading of an initial draft, as well as R. Kakarala also for discussions and sending a copy of [17].

⁵Strictly speaking, \mathcal{F}_{ω} orthornormally projects onto the (coefficient space) of the complement of its kernel.

A [Appendix] Proofs of Lemmas 4.1, 4.2 and 4.3, appearing in Section 4

Proof. [Proof of Lemma 4.1] The set $A^{-1}(\mathcal{B}_m(\epsilon)) \cap \mathcal{B}_r(\delta)$ consists of points in \mathbb{R}^r with 2-norm at most δ , that get mapped to points in \mathbb{R}^m with 2-norm at most ϵ . Since A is surjective with smallest singular value $\sigma > 0$, this set of points is contained in a cylindrical subset of \mathbb{R}^r , with base dimension m , and base radius ϵ/σ , see [24]. The volume of this cylindrical subset is at most $(\epsilon/\sigma)^m \delta^{r-m} \cdot \text{vol}(\mathcal{B}_m(1)) \cdot \text{vol}(\mathcal{B}_{r-m}(1))$, recall we assumed $m \leq r$. On the other hand $\text{vol}(\mathcal{B}_r(\delta)) = \delta^r \cdot \text{vol}(\mathcal{B}_r(1))$. Using these two facts and also the fact that the ℓ -dimensional volume $\text{vol}(\mathcal{B}_\ell(1)) = \pi^{\ell/2}/(\ell/2)!$, we conclude (4.13). \square

Proof. [Proof of Lemma 4.2] As we consider β with bounded 2-norm, it suffices to replace \mathbb{R}^r with $\mathcal{B}_r(\mathbf{0}, \delta)$ for any $\delta > 0$, i.e., it suffices to restrict $\|\beta\|_2 \leq \delta$, for some δ specified in the sequel.

For any bounded β , by assumption ρ_β is Lipschitz, thus there exists some constant C such that the image of any ϵ -ball $\mathcal{B}_n(\epsilon)$ under ρ_β , is contained by in some $(C\epsilon)$ -ball in \mathbb{R}^n . For $k^* > 0$, consider ϵ^{-k^*} number of n -dimensional ϵ -balls, denoted $\mathcal{B}_n(\mathbf{a}_i, \epsilon)$, with various centers \mathbf{a}_i in \mathcal{V} . If $k^* > k$, we can find ϵ^{-k^*} such balls that cover the set \mathcal{V} of interest.

Now for each $\mathcal{B}_n(\mathbf{a}_i, \epsilon)$ in the covering of \mathcal{V} , the image of $\mathcal{B}_n(\mathbf{a}_i, \epsilon)$ under ρ_β contains $\mathbf{0}$, only if $\|\rho_\beta(\mathbf{a}_i)\|_2 < C\epsilon$ for the constants C and ϵ above. For now, we make the following claim that for any $\mathbf{a} \in \mathbb{R}^n$ and some large enough choice for δ

$$\text{vol}(\{\beta \in \mathcal{B}_r(\delta) : \|\rho_\beta(\mathbf{a})\|_2 < C\epsilon\}) \leq C_1 \epsilon^m \quad (\text{A.1})$$

where C_1 is a positive constant. Then for any $\ell > 0$, by a standard argument⁶, the volume of β where at least $\epsilon^{-\ell}$ of the ϵ^{-k^*} images of $\mathcal{B}_n(\mathbf{a}_i, \epsilon)$ contain $\mathbf{0}$ (under ρ_β), is at most $C_1 \epsilon^{m-k^*+\ell}$. In other words, the preimage $\rho_\beta^{-1}(\mathbf{0})$ can be covered by less than $\epsilon^{-\ell}$ number of ϵ -balls, with an exception of maps ρ_β for which the volume of the corresponding β can be made small if $\ell > k^* - m$ and ϵ is small. Thus we conclude when $\ell > k^* - m$ and ϵ goes to 0, we have $\text{boxdim}(\rho_\beta^{-1}(\mathbf{0})) \leq \ell$ for almost every β in $\mathcal{B}_r(\mathbf{0}, \delta)$. As this holds for all $\ell > k^* - m$, and that k^* can be made arbitrarily close to k for sufficiently small ϵ , see [5, 24], we have $\text{boxdim}(\rho_\beta^{-1}(\mathbf{0})) \leq k - m$.

We finish the proof by showing the earlier claim (A.1). Associate $\rho_\beta(\mathbf{a})$ with a linear map A as described in the lemma statement, whereby we assumed that A is surjective. Hence, the positive constant σ as given in the statement of Lemma 4.1 will exist. We then can apply (4.13), by observing that the volume on the LHS of (A.1), equals the volume $\text{vol}(\rho^{-1}(\mathcal{B}_m(C\epsilon)) \cap \mathcal{B}_r(\delta))$ similar to that the LHS of (4.13) (with ϵ replaced by $C\epsilon$). Thus for a large enough choice for δ (where $C/(\sigma\delta) \leq 1$), we can find a constant C_1 that satisfies (A.1). \square

Proof. [Proof of Lemma 4.3] Express $\|\mathbf{Ax}\|_2^2 = \mathbf{x}^T(\mathbf{A}^T\mathbf{A})\mathbf{x} = |\langle \mathbf{A}_i, \mathbf{x} \rangle|^2$, where \mathbf{A}_i equals the i -th row of matrix \mathbf{A} . Call $Z_i = |\langle \mathbf{A}_i, \mathbf{x} \rangle|^2$, and observe $\mathbb{E}Z_1 = \mathbb{E}Z_i = \|\mathbf{x}\|_2^2$, whereby without loss of generality we assume $\|\mathbf{x}\|_2^2 = 1$. We thus want to upper bound the probability $\Pr\{|\sum_{i=1}^n Z_i - m| > m\epsilon\}$. We will only consider one side $\Pr\{\sum_{i=1}^n Z_i - m > m\epsilon\}$, the other side $\Pr\{\sum_{i=1}^n (-Z_i) + m > m\epsilon\}$ can be considered similarly.

By assumption \mathbf{A} has independent rows, the RV's Z_i are mutually independent. Then by Markov's inequality, for any $\theta > 0$

$$\Pr\left\{\sum_{i=1}^n Z_i - m > m\epsilon\right\} \leq e^{-m\theta(\epsilon+1)} \cdot \left(\mathbb{E}e^{\theta Z_1}\right)^m, \quad (\text{A.2})$$

where we used the fact that Z_i 's are identically distributed. Using the fact that the entries of \mathbf{A} are standard normal RV's, then Z_1 is *chi-squared* and for $\theta < 1/2$, and its a standard result that $\mathbb{E}e^{\theta Z_1} = (1 - 2\theta)^{-m/2}$. Substituting this form for $\mathbb{E}e^{\theta Z_1}$ in (A.2), we optimize the upper bound over

⁶For n events $\mathcal{E}_1, \dots, \mathcal{E}_n$, we have that the union bound $\sum_{i=1}^n \Pr\{\mathcal{E}_i\}$ equals $\sum_{i=1}^n \Pr\{\text{at least } i \text{ events } \mathcal{E}_i\}$, see [23], thus we conclude that the union bound is greater than $j \cdot \Pr\{\text{at least } j \text{ events } \mathcal{E}_i\}$ for any $j, 1 \leq j \leq n$.

θ , which requires $\theta = \epsilon/(2 + 2\epsilon) < 1/2$. It follows that the LHS probability of (A.2) is at most $[(1 + \epsilon)e^{-\epsilon}]^{m/2}$, and what we wanted to show follows from the bound $1 + \epsilon \leq \exp(\epsilon - (\epsilon^2 - \epsilon^3)/2)$. \square

References

- [1] ABSIL, P. A., MAHONY, R., AND SEPULCHRE, R. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, 2008.
- [2] ACHLIOPTAS, D. Database-friendly random projections: Johnson-Lindenstrauss with binary coins. *Journal of Computer and System Sciences* 66, 4 (June 2003), 671–687.
- [3] BARANIUK, R., DAVENPORT, M., DEVORE, R., AND WAKIN, M. A Simple Proof of the Restricted Isometry Property for Random Matrices. *Constructive Approximation* 28, 3 (Jan. 2008), 253–263.
- [4] BARANIUK, R. G., AND WAKIN, M. Random Projections of Smooth Manifolds. *Foundations of Computational Mathematics* 9, 1 (2009), 51–77.
- [5] BLUMENSATH, T., AND DAVIES, M. E. Sampling Theorems for Signals From the Union of Finite-Dimensional Linear Subspaces. *IEEE Transactions on Information Theory* 55, 4 (Apr. 2009), 1872–1882.
- [6] CANDÈS, E., AND TAO, T. Near Optimal Signal Recovery From Random Projections : Universal Encoding Strategies? *IEEE Trans. Inform. Theory* 52, 12 (Dec. 2006), 5406–5425.
- [7] CHANDRASEKARAN, V., PARRILO, P. A., AND WILLSKY, A. S. Convex Graph Invariants. *Online: <http://arxiv.org/abs/1012.0623>* (Dec. 2010).
- [8] CHEVALLEY, C. *Theory of Lie Groups I*, first ed. Princeton University Press, 1946.
- [9] CHOI, Y., AND SZPANKOWSKI, W. Compression of Graphical Structures: Fundamental Limits, Algorithms, and Experiments. *IEEE Transactions on Information Theory* 58, 2 (Feb. 2012), 620 – 638.
- [10] CLARKSON, K. L. Tighter bounds for random projections of manifolds. In *24th Annual Symposium on Computational geometry* (2008), pp. 39–48.
- [11] COMON, P., GOLUB, G., LIM, L. H., AND MOURRAIN, B. Symmetric tensors and symmetric tensor rank. *SIAM Journal on Matrix Analysis and Applications* 30, 3 (Sept. 2008), 1254–1279.
- [12] COX, D., LITTLE, J., AND O’SHEA, D. *Ideals, Varieties, and Algorithms*, third ed. Springer, New York, 2007.
- [13] DIACONIS, P. *Group representations in probability and statistics*. Institute of Mathematical Statistics, Lecture Notes–Monograph Series, Vol. 11, 1988.
- [14] DUFRESNE, E. S. *Separating Invariants*. PhD thesis, Queens University, 2008.
- [15] FARIAS, V. F., JAGABATHULA, S., AND SHAH, D. A Nonparametric Approach to Modeling Choice with Limited Data. *Online: <http://arxiv.org/abs/0910.0063>* (2011).
- [16] HUANG, J. *Probabilistic Reasoning and Learning on Permutations: Exploiting Structural Decompositions of the Symmetric Group*. PhD thesis, Carnegie Mellon University, 2011.
- [17] KAKARALA, R. *Triple correlation on groups*. PhD thesis, UC Irvine, 1992.
- [18] KONDOR, R. *Group theoretical methods in machine learning*. PhD thesis, Columbia University, 2008.
- [19] KONDOR, R., SHERVASHIDZE, N., AND BORGWARDT, K. The graphlet spectrum. *Proceedings of the 26th Annual International Conference on Machine Learning (ICML)*, 3 (2009), 1–8.
- [20] LU, Y. M., AND DO, M. N. A Theory for Sampling Signals From a Union of Subspaces. *IEEE Transactions on Signal Processing* 56, 6 (June 2008), 2334 – 2345.
- [21] REZNIK, Y. Coding of Sets of Words. In *Data Compression Conference* (2011), pp. 43 – 52.
- [22] RUDELSON, M., AND VERSHYNIN, R. Non-asymptotic theory of random matrices : extreme singular values. In *Proceedings of the International Congress of Mathematicians* (New Delhi, 2010), Hindustan Book Agency, pp. 1576–1602.
- [23] SATHE, Y. S., PRADHAN, M., AND SHAH, S. P. Inequalities for the Probability of the Occurrence of at least m out of n Events. *Applied Probability* 17, 4 (2012), 1127–1132.
- [24] SAUER, T., YORKE, J. A., AND CASDAGLI, M. Embedology. *Journal of Statistical Physics* 65, 3–4 (1991), 579–616.
- [25] SILBERSTEIN, T. C., SCARABOTTI, F., AND TOLLI, F. *Harmonic Analysis on Finite Groups*. Cambridge University Press, 2008.

- [26] VARSHNEY, L. R., AND GOYAL, V. K. Toward a Source Coding Theory for Sets. In *Data Compression Conference* (2006), pp. 13–22.
- [27] VERSHYNIN, R. Introduction to the non-asymptotic analysis of random matrices. In *Compressed Sensing, Theory and Applications*, Y. Eldar and G. Kutyniok, Eds. Cambridge University Press, 2012, ch. 5, pp. 210–268.
- [28] WOOD, J. Invariant pattern recognition: a review. *Pattern recognition* 29, 1 (1996), 1–17.

SM-I [Supplementary Material] Background on Invariant theory

SM-I.1 The invariant ring always satisfies the discriminability hypothesis: We expect most readers to be unfamiliar with invariant theory. For their convenience, this first set of supplementary material briefly covers results/facts cited and alluded to in the main text. We begin with the connection of invariant theory to *algebraic geometry* - the study of polynomial functions/equations. We discuss the *invariant ring*, *i.e.*, the ring of invariant polynomial functions. We clarify how the \mathcal{G} -invariant \mathcal{F}_ω in (2.4) actually relates to such functions, hence the kernel of \mathcal{F}_ω relates to *algebraic varieties*. We state a result on *seperating invariants* from Defrusne’s thesis (Theorem SM-I.1), that for finite groups the invariant ring has *absolute* discriminative power. We state the results that how the set of canonical points has an manifold structure as an algebraic variety (Theorem SM-I.2). For a good reference text see Cox-Little-O’Shea [12].

We assume some basic *ring theory*. Denote $\mathbb{R}[Z_1, \dots, Z_n]$ to be the ring of n -variate polynomials over \mathbb{R} . For $f \in \mathbb{R}[Z_1, \dots, Z_n]$, let f denote an n -variate polynomial with real coefficients. We think of f as a polynomial *function* with domain \mathbb{R}^n , by letting $f(a_1, \dots, a_n)$ be the evaluation of f at point $(a_1, \dots, a_n) \in \mathbb{R}^n$. By the identification of $\mathbb{R}[\mathcal{X}]$ with \mathbb{R}^n , we also think of f as a function on $\mathbb{R}[\mathcal{X}]$, for some \mathcal{G} -space \mathcal{X} where $\#\mathcal{X} = n$. For some $\mathbf{a} \in \mathbb{R}[\mathcal{X}]$, we write the evaluation as $f(\mathbf{a})$.

Going back to (2.3), we identify $f_{\Omega_{\mathcal{G}}(\mathcal{X}^{\times \omega})}$ with polynomial functions in $\mathbb{R}[Z_1, \dots, Z_n]$, as follows. There exists some $f \in \mathbb{R}[Z_1, \dots, Z_n]$, such that $f_{\Omega_{\mathcal{G}}(\mathcal{X}^{\times \omega})}(\mathbf{a}^{\otimes \omega}) = f(\mathbf{a})$ for any ω -th tensor powers $\mathbf{a}^{\otimes \omega}$, *i.e.*, if the domain $\mathbb{R}[\mathcal{X}^{\times \omega}]$ of the former function is restricted to tensor powers, then the the former function is essentially a polynomial function. This polynomial f that corresponds to $f_{\Omega_{\mathcal{G}}(\mathcal{X}^{\times \omega})}$ must be *homogenous*, *i.e.*, all monomials of f must all be of degree ω .

By the above association of \mathcal{G} -invariants $f_{\Omega_{\mathcal{G}}(\mathcal{X}^{\times \omega})}$ and polynomials f , such an f is a \mathcal{G} -invariant. We formalize the permutation action⁷ of \mathcal{G} on the polynomial ring $\mathbb{R}[Z_1, \dots, Z_n]$. Allow \mathcal{G} to permute the variates Z_i ’s by the identification between \mathbb{R}^n and $\mathbb{R}[\mathcal{X}]$. More specifically for any $g \in \mathcal{G}$, if f^g denotes the polynomial after permuting the variates of f , then for any evaluation under $\mathbf{a} \in \mathbb{R}[\mathcal{X}]$ we have $(f^g)(\mathbf{a}) = f(\mathbf{a}^g)$. Hence if the polynomial f is a \mathcal{G} -invariant, then f must satisfy $f^g = f$ for all $g \in \mathcal{G}$. Invariant theory is the study of the set $\mathbb{R}[Z_1, \dots, Z_n]^{\mathcal{G}}$ of all \mathcal{G} -invariant polynomials, for some group \mathcal{G} . This set is called an **invariant ring** (of \mathcal{G}). Now with reference to the previously discussed polynomial ring $\mathbb{R}[Z_1, \dots, Z_n]$, note that $\mathbb{R}[Z_1, \dots, Z_n]^{\mathcal{G}}$ is a subring of $\mathbb{R}[Z_1, \dots, Z_n]$, and that $\mathbb{R}[Z_1, \dots, Z_n]^{\mathcal{G}}$ contains the constant polynomials. Also $\mathbb{R}[Z_1, \dots, Z_n]^{\mathcal{G}}$ is said to be *graded*, whereby each grade refers to the set of all \mathcal{G} -invariant homogeneous polynomials of a certain degree $\omega \geq 0$, see [12], p. 331. We refer to this set of degree- ω homogeneous polynomials as the ω -**th component** of $\mathbb{R}[Z_1, \dots, Z_n]^{\mathcal{G}}$. Clearly, each ω -th component is closed under \mathbb{R} -linear combinations. In fact, it is known that each such component can be *generated* by κ_ω polynomials $f_1, \dots, f_{\kappa_\omega}$, each f_i corresponding to the i -th orbit invariant $f_{\Omega_{\mathcal{G}, i}}$, recall (2.4). It now becomes clear how the \mathcal{G} -invariant \mathcal{F}_ω corresponds to the ω -th component; each “row” of \mathcal{F}_ω corresponds to a (polynomial) generator. The number κ_ω of generators is computable⁸ by the same equation (2.6).

⁷For simplicity we still focus on permutation actions, though the invariant theoretic results discussed here holds for matrix groups in general.

⁸For matrix groups, we have a more general formula based on Molien’s Theorem [12], p. 340.

At this point one realizes that Algorithm 2.1 in Subsection 2.2 proposes to only use one ω -th component. Evaluating \mathcal{F}_ω only requires polynomial complexity (n^ω operations). But what about the discriminability hypothesis? In the next Supplementary Material SM-II, we explain the connection between each \mathcal{F}_ω and the so-called *multi-correlations* (related to pattern recognition). In particular for the special case $\omega = 3$, Kakarala has applied representation theoretic methods to obtain so-called *completeness* results, or in other words a characterization of the discriminability hypothesis under certain conditions. On the other hand if one is willing to consider the entire invariant ring, the discriminability hypothesis is known to *unconditionally* satisfy for *any* subset in $\mathbb{R}[\mathcal{X}]$. We cite the following result in Dufresne’s thesis, stated here slightly differently⁹.

THEOREM SM-I.1. (COROLLARY 3.2.12, [14], P. 26) *Let \mathcal{G} be a finite group. Let \mathcal{X} be a finite \mathcal{G} -space. Then all ω -th components of the corresponding invariant ring, for all $\omega \leq \#\mathcal{G}$, will be discriminable over the whole data space $\mathbb{R}[\mathcal{X}]$. That is for any fundamental region \mathcal{R} , for any canonical points $\mathbf{a}_1, \mathbf{a}_2 \in \mathbb{R}[\mathcal{X}]_{\mathcal{R}}$, $\mathbf{a}_1 \neq \mathbf{a}_2$, there exists some \mathcal{G} -invariant f in $\mathbb{R}[Z_1, \dots, Z_n]^{\mathcal{G}}$ with degree at most $\#\mathcal{G}$, such that $f(\mathbf{a}_1) \neq f(\mathbf{a}_2)$.*

Recall each \mathcal{F}_ω corresponds to the ω -th component. Hence if all \mathcal{G} -invariants \mathcal{F}_ω , for all $\omega \leq \#\mathcal{G}$, are appropriately made to form a single \mathcal{G} -invariant, then such a \mathcal{G} -invariant will be discriminable over any data set \mathcal{V} . This leads to the following important observation.

FACT SM-I.1. ***The discriminability hypothesis can always be satisfied with large enough computational complexity:** There exists a single \mathcal{G} -invariant corresponding to ω -components, $\omega \leq \#\mathcal{G}$, that for any data set $\mathcal{V} \subset \mathbb{R}[\mathcal{X}]$, satisfies the discriminability hypothesis in both our Whitney embedding Theorem 3.1 and Johnson-Lindenstrauss Theorem 3.2.*

This implies that any bounded, non-sequential data set \mathcal{V} can be appropriately embedded with embedding dimension m tied only to its relevant size k .

However, such an invariant requires $\mathcal{O}(n^{\#\mathcal{G}})$ complexity to compute, exponential in the size of \mathcal{G} - clearly infeasible in practice for most group sizes.

It is not yet known if the size requirements on ω in Theorem SM-I.1 is necessary (in certain cases they can be improved). Now since the same theorem holds for all of $\mathbb{R}[\mathcal{X}]$, one meaningful approach would be relax this requirement, and only consider *specific* subsets of $\mathbb{R}[\mathcal{X}]$. Kakarala adopts a similar strategy for triple-correlations, by obtaining completeness results under certain assumed data conditions (see second set of supplementary material).

SM-I.2 The set of canonical points includes a manifold structure: Another beautiful aspect of invariant theory, is due to its connection with *algebraic geometry*. In particular, there is a remarkable explanation how the set of all canonical points has a manifold-like structure, in the form of an *affine algebraic variety* [12], pp. 345-353. An (affine) algebraic variety is a set of points, whereby there exists a set of polynomial equations, for which is satisfied by every point in this set. For example, the kernel of the \mathcal{G} -invariant \mathcal{F}_ω in (2.4) is related to the following algebraic variety

$$\{\mathbf{a} \in \mathbb{R}[\mathcal{X}] : f_i(\mathbf{a}) = 0, 1 \leq i \leq \kappa_\omega\}, \quad (\text{SM-I.1})$$

where $f_i(\mathbf{a}) = f_{\Omega_{\mathcal{G},i}(\mathcal{X}^{\times \omega})}(\mathbf{a}^{\otimes \omega})$. For the same polynomials f_i , the following set is also an algebraic variety

$$\{(\mathbf{a}_1, \mathbf{a}_2) \in \mathbb{R}[\mathcal{X}] \times \mathbb{R}[\mathcal{X}] : f_i(\mathbf{a}_1) - f_i(\mathbf{a}_2) = 0, 1 \leq i \leq \kappa_\omega\}, \quad (\text{SM-I.2})$$

whereby this second set (SM-I.2) contains pairs of points in $\mathbb{R}[\mathcal{X}]$ that *cannot* be discriminated by the \mathcal{G} -invariant \mathcal{F}_ω . In theory, the set could be computed by *elimination theory* using a *Gröbner*

⁹The statement in [14] uses a stronger notion of discriminability, called a *geometric separating set*, see Definition 3.2.1, p. 15. Also it holds for general matrix groups.

basis, see [12], ch. 3, which will obtaining useful characterizations of such pairs of points $(\mathbf{a}_1, \mathbf{a}_2)$. Though such an approach can be unwieldy for large n , it does suggest a possible algebraic geometry view of characterizing discriminability of invariants, besides the representation theoretic techniques of Kakarala's. Also Kakarala's techniques currently only hold for triple-correlations (*i.e.*, $\omega = 3$), whereas here ω could be arbitrary.

The algebraic variety structure of the set of canonical points is a little more complicated to explain, and requires the *algebraic closure* of \mathbb{R} to the complex field \mathbb{C} . Take a generating set of the invariant ring $\mathbb{C}[Z_1, \dots, Z_n]^{\mathcal{G}}$ over \mathbb{C} , say f_1, \dots, f_ℓ for some $\ell \geq 1$, and form a map $\rho : \mathbb{C}[\mathcal{X}] \rightarrow \mathbb{C}^\ell : \mathbf{a} \mapsto (f_1(\mathbf{a}), \dots, f_\ell(\mathbf{a}))$, where $\mathbb{C}[\mathcal{X}]$ is the *complexification* of $\mathbb{R}[\mathcal{X}]$. Recall the notation $\mathbb{C}[\mathcal{X}]_{\mathcal{R}}$, which means a set of canonical points in $\mathbb{C}[\mathcal{X}]$ lying in some fundamental region \mathcal{R} . There exists an invariant theoretic result that says that $\mathbb{C}[\mathcal{X}]_{\mathcal{R}}$ is in bijection with the *image* of ρ , whereby this image is actually an algebraic variety. The set of polynomial equations that describe the image comes from the generators of a special *ideal* of the ring $\mathbb{C}[Y_1, \dots, Y_\ell]$ of ℓ -variate polynomials, where ℓ is the number of generators f_i of the invariant ring. This ideal, known as the **ideal of relations**, contain all β in $\mathbb{C}[Y_1, \dots, Y_\ell]$ whereby $\beta(f_1, \dots, f_\ell)$ is identically zero; here $\beta(f_1, \dots, f_\ell)$ is thought of as a polynomial in the variates Z_i 's. This result is stated as follows.

THEOREM SM-I.2. (THEOREM 10, [12], P. 351) *Let f_1, \dots, f_ℓ generate the invariant ring $\mathbb{C}[Z_1, \dots, Z_n]^{\mathcal{G}}$, for some $\ell \geq 1$. Let $\rho : \mathbb{C}[\mathcal{X}] \rightarrow \mathbb{C}^\ell : \mathbf{a} \mapsto (f_1(\mathbf{a}), \dots, f_\ell(\mathbf{a}))$.*

Let β_1, \dots, β_r generate the ideal of relations in the ring $\mathbb{C}[Y_1, \dots, Y_\ell]$, for some $r \geq 1$. Consider the algebraic variety

$$\{(b_1, \dots, b_r) \in \mathbb{C}^r : \beta_i(b_1, \dots, b_r) = 0, 1 \leq i \leq r\} \quad (\text{SM-I.3})$$

Then the image of ρ is surjective over the algebraic variety (SM-I.3). In fact if we restrict ρ over the domain $\mathbb{C}[\mathcal{X}]_{\mathcal{R}}$ for any fundamental region \mathcal{R} , then ρ with this restriction of domain, becomes bijective.

Theorem SM-I.2 remarkably shows how the set of canonical points, after passing through this map ρ , has the manifold structure of the algebraic variety (SM-I.3). This brings to mind the possibility of applying *manifold learning* techniques to learn the canonical points. However until one derives an analogue of Theorem SM-I.2 for the reals, one needs to work in \mathbb{C} .

SM-II [Supplementary Material] Completeness results for triple-correlation

SM-II.1 Multi-correlations are connected with invariant theory: Auto- and triple-correlation functions have been employed as invariants in pattern recognition [17–19], though the presentation has always been disparate from invariant theory. The first goal of this second set of supplementary material, is to provide unification. We begin by clarifying how a generalization of such functions (that we call *multi-correlations*) are one and the same to the graded components of the invariant ring (see previous Supplementary Material SM-I). Then next, for the sake of most readers not familiar with Kakarala's completeness results for the triple-correlation, we provide a primer in Subsection SM-II.2).

For correlation functions studied pattern recognition, the group action is limited to *transitive* permutation actions. Recall the two examples given in Subsection 2.1. For this special case, the \mathcal{G} -space \mathcal{X} is referred to as a **homogeneous space**. To explain correlations, we require the following notion of \mathcal{G} itself as a homogeneous space.

Example. [\mathcal{G} as a homogeneous space]: For an abstract group \mathcal{G} , define a action of \mathcal{G} on itself, where for any $g \in \mathcal{G}$, we have the image $g(\sigma) = g\sigma$ for any $\sigma \in \mathcal{G}$, *i.e.*, \mathcal{G} acts on itself by left multiplication. This is a transitive action, so \mathcal{G} (as a set) is a homogeneous space.

The last example admits discussion of the vector space $\mathbb{R}[\mathcal{G}]$; we consider \mathcal{G} as the set \mathcal{X} . Let \mathbf{z} denote an element in $\mathbb{R}[\mathcal{G}]$, where z_g denotes an indexed element of \mathbf{z} for $g \in \mathcal{G}$. For any $\mathbf{z} \in \mathbb{R}[\mathcal{G}]$, the **multi-correlation** $\mathcal{A}_{\mathbf{z}}^{(\omega)}$ for some $\omega \geq 1$, is given as

$$\mathcal{A}_{\mathbf{z}}^{(\omega)}(g_1, \dots, g_{\omega-1}) = \sum_{\sigma \in \mathcal{G}} z_{\sigma} z_{\sigma g_1} \cdots z_{\sigma g_{\omega-1}}, \quad (\text{SM-II.4})$$

where for j , $1 \leq j < \omega$ we have $g_j \in \mathcal{G}$. The cases $\omega = 2$ and $\omega = 3$ specialize respectively to the auto- and triple-correlations. For any $\omega \geq 1$, the function $\mathcal{A}_{\mathbf{z}}^{(\omega)}$ is a \mathcal{G} -invariant, *i.e.*, for any $\alpha \in \mathcal{G}$, we have $\mathcal{A}_{\mathbf{z}^{\alpha}}^{(\omega)} = \mathcal{A}_{\mathbf{z}}^{(\omega)}$; to verify this, simply evaluate (SM-II.4) with \mathbf{z}^{α} and put $(\mathbf{z}^{\alpha})_{\sigma} = z_{\alpha^{-1}\sigma}$ for any $\sigma \in \mathcal{G}$.

While the (correlation) functions (SM-II.4) seem to be only defined for the space $\mathbb{R}[\mathcal{G}]$, we can accommodate any \mathcal{G} -space \mathcal{X} , by *extending* elements in $\mathbb{R}[\mathcal{X}]$ to $\mathbb{R}[\mathcal{G}]$. Let x_1 denote an element in \mathcal{X} that has been (arbitrarily) chosen and fixed. Using this x_1 then for any $\mathbf{a} \in \mathbb{R}[\mathcal{X}]$, the **extension** of \mathbf{a} , denoted $\bar{\mathbf{a}}$, satisfies

$$\bar{a}_g = a_{g(x_1)}, \quad \text{for all } g \in \mathcal{G}. \quad (\text{SM-II.5})$$

The **stabilizer** of the fixed element x_1 , denoted \mathcal{S}_{x_1} , is the set of group elements in \mathcal{G} that leave x_1 un-moved, *i.e.*, $\mathcal{S}_{x_1} = \{g \in \mathcal{G} : g(x_1) = x_1\}$. Clearly \mathcal{S}_{x_1} will be a subgroup of \mathcal{G} . Since we do not discuss other stabilizer subgroups in the sequel, we will drop the subscript x_1 from \mathcal{S}_{x_1} and simply write \mathcal{S} throughout. The relationship (SM-II.5) relates \mathcal{S} to extensions of vectors in $\mathbb{R}[\mathcal{X}]$, whereby note that any extension $\bar{\mathbf{a}}$ is *constant over left-cosets* of \mathcal{S} in \mathcal{G} , *i.e.*, for any $g \in \mathcal{G}$, we have $\bar{a}_{gs} = \bar{a}_g$ for any $s \in \mathcal{S}$. Hence when considering homogeneous spaces \mathcal{X} we only need to evaluate (SM-II.4) (for $\mathcal{A}_{\bar{\mathbf{a}}}^{(\omega)}$ where $\mathbf{a} \in \mathbb{R}[\mathcal{X}]$) at points $\{(t_{i_1}, \dots, t_{i_{\omega-1}}) : 1 \leq i_1, \dots, i_{\omega-1} \leq n\}$, where each t_j is a **left-coset representative**. There are at most $n^{\omega-1}$ such points, where $n = \#\mathcal{X}$. For the previously fixed x_1 , enumerate the rest of the elements in \mathcal{X} as x_2, x_3, \dots, x_n , and fix t_j to send x_1 to x_j (possible only when \mathcal{G} acts transitively on \mathcal{X}). Note $n = \#\mathcal{X} = \#\mathcal{G}/\#\mathcal{S}$. To conclude, extensions allow us to synonymously discuss correlations for $\mathbb{R}[\mathcal{G}]$, and $\mathbb{R}[\mathcal{X}]$ for any homogeneous \mathcal{G} -space \mathcal{X} .

We proceed to show how the multi-correlation (SM-II.4) for some $\omega \geq 1$, is related to the ω -th component of the invariant ring. We do this by specifying the connection with \mathcal{G} -invariant \mathcal{F}_{ω} in (2.4), which was already established to “generate” the ω -th degree polynomials in the ring. For any $\mathbf{a} \in \mathbb{R}[\mathcal{X}]$, we calculate the multi-correlation $\mathcal{A}_{\bar{\mathbf{a}}}^{(\omega)}$ as follows

$$\begin{aligned} \mathcal{A}_{\bar{\mathbf{a}}}^{(\omega)}(t_{i_1}, \dots, t_{i_{\omega-1}}) &= \sum_{\sigma \in \mathcal{G}} \bar{a}_{\sigma} \bar{a}_{\sigma t_{i_1}} \cdots \bar{a}_{\sigma t_{i_{\omega-1}}} \\ &\stackrel{(a)}{=} \sum_{j=1}^n \sum_{s \in \mathcal{S}} \bar{a}_{t_j s} \bar{a}_{t_j s t_{i_1}} \cdots \bar{a}_{t_j s t_{i_{\omega-1}}} \\ &\stackrel{(b)}{=} \sum_{j=1}^n \sum_{s \in \mathcal{S}} a_{x_j} a_{(t_j s t_{i_1})(x_1)} \cdots a_{(t_j s t_{i_{\omega-1}})(x_1)} \\ &= \sum_{j=1}^n a_{x_j} \sum_{s \in \mathcal{S}} a_{(t_j s)(x_{i_1})} \cdots a_{(t_j s)(x_{i_{\omega-1}})} \end{aligned} \quad (\text{SM-II.6})$$

where in (a) we apply $\sigma = t_j s$ for some t_j , in (b) we apply (SM-II.5) and $\bar{a}_{t_j s} = a_{(t_j s)(x_1)} = a_{t_j(x_1)} = a_{x_j}$, and the last equality follows by definition $t_j(x_1) = x_j$. We notice the following from the final expression (SM-II.6). For each j , $1 \leq j \leq n$, the second summation really runs over indexes over $\mathcal{X}^{\times(\omega-1)}$ in the set $\{t_j(\mathbf{x}^{(1:\omega-1)}) : \mathbf{x}^{(1:\omega-1)} \in \Omega_{\mathcal{S}}(\mathcal{X}^{\times(\omega-1)})\}$, where $\Omega_{\mathcal{S}}(\mathcal{X}^{\times(\omega-1)})$ is the \mathcal{S} -orbit (over $\mathcal{X}^{\times(\omega-1)}$) that contains $(x_{i_1}, \dots, x_{i_{\omega-1}})$. The LHS and RHS of (SM-II.6) are really determined by

the indices $i_1, \dots, i_{\omega-1}$, for at most $n^{\omega-1}$ such choices.

We notice the following connection between the final expression in (SM-II.6) and the \mathcal{G} -invariant as applied in Algorithm 2.1. First, there is a one-to-one correspondence between \mathcal{G} -orbits on $\mathcal{X}^{\times\omega}$, and \mathcal{S} -orbits on $\mathcal{X}^{\times(\omega-1)}$. This correspondence is obtained for $\Omega_{\mathcal{G}}(\mathcal{X}^{\times\omega})$, by identifying $\Omega_{\mathcal{S}}(\mathcal{X}^{\times(\omega-1)})$ with the subset $\{\mathbf{x}^{(1:\omega-1)} : (\mathbf{x}^{(1:\omega-1)}, x_1) \in \Omega\}$ of $\mathcal{X}^{\times(\omega-1)}$. Secondly for any \mathcal{G} -orbit $\Omega = \Omega_{\mathcal{G}}(\mathcal{X}^{\times\omega})$ on $\mathcal{X}^{\times\omega}$, by the corresponding ω -array $\llbracket b_{\mathbf{x}^{(1:\omega)}} \rrbracket$ in (2.1), we can express (see (2.4))

$$f_{\Omega}(\mathbf{a}^{\otimes\omega}) = \sum_{j=1}^n a_{x_j} \left(\sum_{(x^{(1)}, \dots, x^{(\omega-1)}) \in \Omega'_j} a_{x^{(1)}} \cdots a_{x^{(\omega-1)}} \right)$$

where for each j , $1 \leq j \leq n$, we have $\Omega'_j = \{\mathbf{x}^{(1:\omega-1)} : (\mathbf{x}^{(1:\omega-1)}, x_j) \in \Omega\}$. Note that Ω'_j is simply an orbit of the subgroup $t_j \mathcal{S} t_j^{-1}$ that stabilizes x_j , whereby $\Omega'_j = \Omega_{\mathcal{S}}(\mathcal{X}^{\times(\omega-1)})$, the \mathcal{S} -orbit previously identified with the \mathcal{G} -orbit Ω . Recall from the proof of Proposition 2.1 that the $(t_j \mathcal{S} t_j^{-1})$ -orbit is simply the set $\{t_j(\mathbf{x}^{(1:\omega-1)}) : \mathbf{x}^{(1:\omega-1)} \in \mathcal{X}^{\times(\omega-1)}\}$. Finally, compare with (SM-II.6) by taking $(x_{i_1}, \dots, x_{i_{\omega-1}}) \in \Omega_{\mathcal{S}}(\mathcal{X}^{\times(\omega-1)})$ (determined by the indices $i_1, \dots, i_{\omega-1}$), and conclude the following result.

PROPOSITION SM-II.1. *Let $\mathbf{a} \in \mathbb{R}[\mathcal{X}]$. Let $\Omega_{\mathcal{S},1}(\mathcal{X}^{\times(\omega-1)}), \dots, \Omega_{\mathcal{S},\kappa_{\omega}}(\mathcal{X}^{\times(\omega-1)})$ denote the κ_{ω} number of \mathcal{S} -orbits on $\mathcal{X}^{\times(\omega-1)}$. Then firstly for an extension $\bar{\mathbf{a}}$, the multi-correlation $\mathcal{A}_{\bar{\mathbf{a}}}^{(\omega)}$ has at most κ_{ω} unique evaluations, found at the points $(t_{i_1}, \dots, t_{i_{\omega-1}})$ corresponding to the representatives $(x_{i_1}, \dots, x_{i_{\omega-1}})$ of the \mathcal{S} -orbits.*

Secondly, the output $\mathcal{F}_{\omega}(\mathbf{a}^{\otimes\omega})$ of Algorithm 2.1 is equivalent to the multi-correlation $\mathcal{A}_{\bar{\mathbf{a}}}^{(\omega)}$ for the extension $\bar{\mathbf{a}}$, whereby evaluation at the point $(t_{i_1}, \dots, t_{i_{\omega-1}})$ corresponding to $(x_{i_1}, \dots, x_{i_{\omega-1}})$, is equal to the value of $f_{\Omega_{\mathcal{G}}(\mathcal{X}^{\times\omega})}(\mathbf{a}^{\otimes\omega})$, see (2.4), where the \mathcal{G} -orbit $\Omega_{\mathcal{G}}(\mathcal{X}^{\times\omega})$ corresponds to the \mathcal{S} -orbit that contains $(x_{i_1}, \dots, x_{i_{\omega-1}})$.

The second part of Proposition SM-II.1 proves the intended equivalence between the \mathcal{G} -invariants in 2.4 and the multi-correlations. This proposition establishes a connection between Kakarala's representation theoretic analysis, discussed in the sequel, and the invariant theory discussed in Supplementary Material SM-I.

SM-II.2 Kakarala's completeness results for triple-correlation: This subsection provides a brief introduction to *representation theoretic* techniques for showing completeness of the triple correlation. We discuss a constructive algorithm for finite cyclic groups (which more generally also applies to finite abelian groups), and Kakarala's completeness result for compact groups. Note that compact groups include finite groups under the discrete topology. Good references to this material include the textbook [25], and Kakarala's and Kondor's theses [17, 18].

Here we let \mathcal{V} denote a finite-dimensional vector space. A **representation** of a group \mathcal{G} over \mathcal{V} , is an action of \mathcal{G} on the vector space \mathcal{V} ; for any $g \in \mathcal{G}$, each $\mathbf{z} \in \mathcal{V}$ is sent to $\rho(g)\mathbf{z}$, whereby any $\rho(g)$ is an invertible linear map. For example suppose $\mathcal{V} = \mathbb{R}[\mathcal{G}]$, and for $g \in \mathcal{G}$ set $\rho(g)$ to be a 0-1 matrix in $\mathbb{R}^{\mathcal{G} \times \mathcal{G}}$ whose h, σ -th element $(\rho(g))_{h,\sigma}$ equals 1 i.f.f. $h = g\sigma$. This representation, called the **left-regular representation**, is in fact related to the previous example of \mathcal{G} acting on itself (i.e., \mathcal{G} is a homogeneous \mathcal{G} -space).

A representation (ρ, \mathcal{V}) is said to be **irreducible**, if the subspace of \mathcal{V} invariant under the representation action, is trivial (i.e., the invariant subspace equals either $\{\mathbf{0}\}$ or \mathcal{V}). An **unitary** representation (ρ, \mathcal{V}) preserves the inner product on \mathcal{V} , i.e., for all $g \in \mathcal{G}$ we have $\langle \rho(g)\mathbf{z}, \rho(g)\mathbf{z}' \rangle = \langle \mathbf{z}, \mathbf{z}' \rangle$ for any $\mathbf{z}, \mathbf{z}' \in \mathcal{V}$. Two representations (ρ_1, \mathcal{V}_1) and (ρ_2, \mathcal{V}_2) are said to be **equivalent**, if there exists a linear bijection $A : \mathcal{V}_2 \rightarrow \mathcal{V}_1$ such that $\rho_1(g)A = A\rho_2$ for all $g \in \mathcal{G}$.

The **dual** of a finite group \mathcal{G} , denoted $\widehat{\mathcal{G}}$, is the complete set of irreducible pairwise non-equivalent (unitary) representations of \mathcal{G} . If \mathcal{G} is finite then so is $\widehat{\mathcal{G}}$. The machinery to obtain $\widehat{\mathcal{G}}$, from the left-regular representation, is given by the *Peter-Weyl theorem* (see [25], pp. 85-86, for the statement for finite \mathcal{G}). The following is the analogue of the Fourier transform, stated for finite \mathcal{G} .

DEFINITION SM-II.1. (c.f., [25], P. 99) Let $\mathbf{z} \in \mathbb{R}[\mathcal{G}]$. Let \mathcal{G} be a finite group with finite dual $\widehat{\mathcal{G}}$. The (abstract) **Fourier transform** component of \mathbf{z} with respect to a irreducible (unitary) representation (ρ, \mathcal{V}) , is the linear operator $\hat{\mathbf{z}}(\rho) : \mathcal{V} \rightarrow \mathcal{V}$ defined by

$$\hat{\mathbf{z}}(\rho) = \sum_{g \in \mathcal{G}} z_g \cdot \rho(g). \quad (\text{SM-II.7})$$

The techniques here will be very related to this Fourier transform. In what follows, we need to consider the product group $\mathcal{G} \times \mathcal{G}$, and its dual $\widehat{\mathcal{G} \times \mathcal{G}}$. Here, each $(\rho, \mathcal{V}) \in \widehat{\mathcal{G} \times \mathcal{G}}$ has maps $\rho(g, h)$ indexed by an element pair $g, h \in \mathcal{G}$. For the the triple correlation $\mathcal{A}_{\mathbf{z}}^{(3)}$ of any $\mathbf{z} \in \mathbb{R}[\mathcal{G}]$, we now elucidate an illuminating structure of a Fourier transform component, specially¹⁰ denoted $\mathcal{B}_{\mathbf{z}}(\rho)$. Consider two elements $\mathbf{z}_1, \mathbf{z}_2 \in \mathbb{R}[\mathcal{G} \times \mathcal{G}]$ related to $\mathbf{z} \in \mathbb{R}[\mathcal{G}]$, as follows. For \mathbf{z}_1 , set $(\mathbf{z}_1)_{(g,g)} = z_g$ for all $g \in \mathcal{G}$ and $(\mathbf{z}_1)_{(g,h)} = 0$ when $h \neq g$. For \mathbf{z}_2 , set $(\mathbf{z}_2)_{(g,h)} = z_g z_h$ for all $g, h \in \mathcal{G}$. Let \dagger denote *complex conjugation*. Then for any $(\rho, \mathcal{V}) \in \widehat{\mathcal{G} \times \mathcal{G}}$, we see that

$$\begin{aligned} (\widehat{\mathbf{z}}_1(\rho))^\dagger \widehat{\mathbf{z}}_2(\rho) &= \left(\sum_{\sigma \in \mathcal{G}} z_\sigma \cdot \rho(\sigma^{-1}, \sigma^{-1}) \right) \cdot \left(\sum_{g, h \in \mathcal{G}} z_g z_h \cdot \rho(g, h) \right) \\ &= \sum_{h, g \in \mathcal{G}} \sum_{\sigma \in \mathcal{G}} z_\sigma z_g z_h \cdot \rho(\sigma^{-1}g, \sigma^{-1}h) \\ &= \sum_{h, g \in \mathcal{G}} \sum_{\sigma \in \mathcal{G}} z_\sigma z_{\sigma g} z_{\sigma h} \cdot \rho(g, h) \\ &= \sum_{g, h \in \mathcal{G}} \mathcal{A}_{\mathbf{z}}^{(3)}(g, h) \cdot \rho(g, h) = \mathcal{B}_{\mathbf{z}}(\rho), \end{aligned} \quad (\text{SM-II.8})$$

where the second last equality follows from the definition (SM-II.4) of the triple correlation $\mathcal{A}_{\mathbf{z}}^{(3)}$.

We proceed to further manipulate the LHS of (SM-II.8). Each (ρ, \mathcal{V}) in $\widehat{\mathcal{G} \times \mathcal{G}}$ can be expressed as $(\rho_1 \otimes \rho_2, \mathcal{V}_1 \otimes \mathcal{V}_2)$, where $(\rho_1, \mathcal{V}_1), (\rho_2, \mathcal{V}_2) \in \widehat{\mathcal{G}}$, where $\rho(g, h) = \rho_1(g) \otimes \rho_2(h)$, see [25], p. 272. Thus for $\widehat{\mathbf{z}}_2(\rho)$ in (SM-II.8), $\rho = \rho_1 \otimes \rho_2$, we conclude

$$\widehat{\mathbf{z}}_2(\rho_1 \otimes \rho_2) = \hat{\mathbf{z}}(\rho_1) \otimes \hat{\mathbf{z}}(\rho_2), \quad (\text{SM-II.9})$$

where the RHS are two Fourier transforms of \mathbf{z} in $\mathbb{R}[\mathcal{X}]$, corresponding to representations $(\rho_1, \mathcal{V}_1), (\rho_2, \mathcal{V}_2) \in \widehat{\mathcal{G}}$. Next we require the notion¹¹ of a *direct sum representation* $(\rho_1 \oplus \rho_2, \mathcal{V}_1 \oplus \mathcal{V}_2)$ of two representations (ρ_1, \mathcal{V}_1) and (ρ_2, \mathcal{V}_2) of \mathcal{G} , where $\mathcal{V}_1, \mathcal{V}_2$ are orthogonal. In the direct sum for all $g \in \mathcal{G}$, we mean that $\varrho_1(g)$ leaves \mathcal{V}_2 invariant, and $\varrho_2(g)$ leaves \mathcal{V}_1 invariant. The tensor product representation $\rho_1 \otimes \rho_2$ can be expressed as direct sums of representations in $\widehat{\mathcal{G}}$, i.e.,

$$\rho_1 \otimes \rho_2 \equiv \bigoplus_{\varrho \in \widehat{\mathcal{G}}} \varrho^{\otimes m_{\rho_1, \rho_2}(\varrho)} \quad (\text{SM-II.10})$$

where \equiv denotes equivalence in representations (under some linear operator $A_{\rho_1, \rho_2} : \mathcal{V} \rightarrow \mathcal{V}'$ where \mathcal{V}' is some subspace of $\mathbb{R}[\mathcal{X}]$), and the notation $\varrho^{\otimes \ell}$ for $\varrho \in \widehat{\mathcal{G}}$, $\ell \in \mathbb{Z}$, means the representation $\varrho \otimes \cdots \otimes \varrho$ formed by ℓ copies of ϱ , and finally $m_{\rho_1, \rho_2} : \widehat{\mathcal{G}} \rightarrow \mathbb{Z}$ returns for each ϱ in $\widehat{\mathcal{G}}$, the number of copies in the tensor product. From (SM-II.10) we can conclude for $\widehat{\mathbf{z}}_1(\rho)$ in (SM-II.8), where

¹⁰The \mathcal{B} stands for *bi-spectrum*, a term for the (2-dimensional) Fourier transform of the triple correlation.

¹¹The direct sum $\mathcal{V}_1 \oplus \mathcal{V}_2$ of vector spaces equals $\{\mathbf{v}_1 + \mathbf{v}_2 : \mathbf{v}_1 \in \mathcal{V}_1, \mathbf{v}_2 \in \mathcal{V}_2\}$.

$$\rho = \rho_1 \otimes \rho_2,$$

$$\widehat{\mathbf{z}}_1(\rho_1 \otimes \rho_2) \equiv \bigoplus_{\varrho \in \widehat{\mathcal{G}}} (\widehat{\mathbf{z}}(\varrho))^{\otimes m_{\rho_1, \rho_2}(\varrho)} \quad (\text{SM-II.11})$$

where \equiv means the same equivalence earlier in (SM-II.10). By the identity $\mathcal{B}_{\mathbf{z}}(\rho) = (\widehat{\mathbf{z}}_1(\rho))^\dagger \widehat{\mathbf{z}}_2(\rho)$ developed in (SM-II.8), we conclude where $\rho = \rho_1 \otimes \rho_2$ the following

$$\mathcal{B}_{\mathbf{z}}(\rho_1 \otimes \rho_2) A_{\rho_1 \otimes \rho_2} = \left[\bigoplus_{\varrho \in \widehat{\mathcal{G}}} (\widehat{\mathbf{z}}(\varrho))^{\otimes m_{\rho_1, \rho_2}(\varrho)} \right]^\dagger A_{\rho_1 \otimes \rho_2} \widehat{\mathbf{z}}(\rho_1) \otimes \widehat{\mathbf{z}}(\rho_2). \quad (\text{SM-II.12})$$

where $A_{\rho_1 \otimes \rho_2}$ makes the equivalence (SM-II.10). From (SM-II.12), we can now describe an algorithm that recovers the Fourier coefficients $\widehat{\mathbf{z}}(\rho)$ from that of the triple-correlation $\mathcal{A}_{\mathbf{z}}^{(3)}$ (*i.e.*, from $\mathcal{B}_{\mathbf{z}}(\rho_1 \otimes \rho_2)$). Then by a *Fourier inversion theorem*, [25], p. 100, we contain obtain from $\widehat{\mathbf{z}}(\rho)$ the data \mathbf{z} .

A condition will be required for the algorithm to work:

$$\text{for all } (\rho, \mathcal{V}) \in \widehat{\mathcal{G}}, \quad \widehat{\mathbf{z}}(\rho) \text{ is an invertible map.} \quad (\text{SM-II.13})$$

If (SM-II.13) holds, then for all $\rho_1, \rho_2 \in \widehat{\mathcal{G}}$ the following quantity

$$\mathcal{B}'_{\mathbf{z}}(\rho_1 \otimes \rho_2) = \mathcal{B}_{\mathbf{z}}(\rho_1 \otimes \rho_2) A_{\rho_1 \otimes \rho_2} \widehat{\mathbf{z}}^{-1}(\rho_1) \otimes \widehat{\mathbf{z}}^{-1}(\rho_2) A_{\rho_1 \otimes \rho_2}^\dagger \quad (\text{SM-II.14})$$

is well-defined, where $A_{\rho_1 \otimes \rho_2}^\dagger$ is the adjoint of $A_{\rho_1 \otimes \rho_2}$ with complex conjugation. Let $(\mathbb{1}, \mathcal{V})$ denote the **trivial representation** whereby $\mathbb{1}(g) = 1$ for all $g \in \mathcal{G}$. We see that

$$\begin{aligned} \widehat{\mathbf{z}}_1(\mathbb{1} \otimes \rho) &= \widehat{\mathbf{z}}(\rho), \\ \widehat{\mathbf{z}}_2(\mathbb{1} \otimes \rho) &= \widehat{\mathbf{z}}(\mathbb{1}) \cdot \widehat{\mathbf{z}}(\rho), \end{aligned} \quad (\text{SM-II.15})$$

which follows from (SM-II.11) and (SM-II.9). Then from (SM-II.12) the following algorithm¹², under the existence of an appropriate labeling $\varrho_1, \varrho_2, \varrho_3, \dots$ of representations in $\widehat{\mathcal{G}}$ (where $\varrho_1 = \mathbb{1}$), will perform the promised task.

ALGORITHM SM-II.1. To obtain Fourier coefficients $\widehat{\mathbf{z}}(\rho)$ from $\mathcal{B}_{\mathbf{z}}(\rho_1 \otimes \rho_2)$, where $\rho, \rho_1, \rho_2 \in \widehat{\mathcal{G}}$

- As $\mathcal{B}_{\mathbf{z}}(\mathbb{1} \otimes \mathbb{1}) = (\widehat{\mathbf{z}}(\mathbb{1}))^3$ holds from (SM-II.8) and (SM-II.15), compute $\widehat{\mathbf{z}}(\mathbb{1}) = \widehat{\mathbf{z}}(\varrho_1)$.
- As $\mathcal{B}_{\mathbf{z}}(\mathbb{1} \otimes \varrho_2) = \widehat{\mathbf{z}}(\mathbb{1}) \cdot (\widehat{\mathbf{z}}(\varrho_2))^\dagger \widehat{\mathbf{z}}(\varrho_2)$ holds from (SM-II.8) and (SM-II.15), compute $\widehat{\mathbf{z}}(\varrho_2)$.
 - Note that since $\mathbf{z}^{\widehat{\alpha}}(\varrho_2) = \rho_2(\alpha) \widehat{\mathbf{z}}(\varrho_2)$ for any $\alpha \in \mathcal{G}$, we can only determine $\widehat{\mathbf{z}}(\varrho_2)$ up to \mathcal{G} -invariance (*i.e.*, if $\widehat{\mathbf{z}}(\varrho_2)$ solves the above expression, then so does $\rho_2(\alpha) \widehat{\mathbf{z}}(\varrho_2)$ for any $\alpha \in \mathcal{G}$).
- For $\varrho_3, \varrho_4, \dots$, use the following iteration derived from both (SM-II.12) and (SM-II.14). For $\ell \geq 3$, use

$$\mathcal{B}'_{\mathbf{z}}(\varrho_{\ell-1} \otimes \varrho_2) = \widehat{\mathbf{z}}(\varrho_\ell)^\dagger \oplus M_{\ell-1}$$

to solve for $\widehat{\mathbf{z}}(\varrho_\ell)$ where the LHS will be known using previous computations. where $M_{\ell-1}$ is the remainder term in the RHS of (SM-II.10) for $\varrho_{\ell-1} \otimes \varrho_2$, after pulling out one copy of ϱ_ℓ .

Now for the final step of Algorithm SM-II.1 to work, the labeling $\varrho_1, \varrho_2, \varrho_3, \dots$ must allow $\widehat{\mathbf{z}}(\varrho_\ell)$ to be pulled out in each ℓ -th step. Unfortunately in general for finite groups \mathcal{G} , this labeling is

¹²This steps of this algorithm was not stated as clearly in previous work, hence it is valuable to record them here.

unknown. On the other hand if \mathcal{G} is cyclic, the representations $(\varrho_\ell, \mathcal{V}) \in \widehat{\mathcal{G}}$, $1 \leq \ell \leq \#\mathcal{G}$, possess a “cyclic group structure”, see [25], p. 274. In particular, there exists some choice for labeling $\varrho_1, \varrho_2, \varrho_3, \dots$, such that we can express for any $2 \leq \ell \leq \#\mathcal{G}$

$$\varrho_\ell = \varrho_{\ell-1} \otimes \varrho_2$$

using some special choice for ϱ_2 . Hence for finite cyclic groups, Algorithm SM-II.1 will work as long condition (SM-II.13) is met. Also for finite abelian groups in general, which are always isomorphic to direct product of a finite number of finite cyclic groups, appropriate extensions can be perused. In conclusion, Algorithm SM-II.1 is a constructive proof of a completeness result (under the above appropriate conditions), that $\mathcal{A}_{\mathbf{z}}^{(3)} = \mathcal{A}_{\mathbf{z}'}^{(3)}$ if and only if \mathbf{z}' must be some obtainable from \mathbf{z} by some $g \in \mathcal{G}$.

Using the condition (SM-II.13), Kakarala proved a remarkable completeness result of the same vein, for the large class of compact groups (which also includes some infinite groups - under appropriate generalization of the vector space $\mathbb{R}[\mathcal{G}]$, the Fourier transform in Definition SM-II.1, and the dual $\widehat{\mathcal{G}}$, see [17] for details).

THEOREM SM-II.1. (*c.f.*, [17]) *Let \mathcal{G} be a compact group, and let $\widehat{\mathcal{G}}$ be its dual. Let \mathbf{z} be any arbitrary function in $\mathbb{R}[\mathcal{G}]$, for which we assume that condition (SM-II.13) is met. Then the triple-correlation $\mathcal{A}_{\mathbf{z}}^{(3)}$ of \mathbf{z} , equals another $\mathcal{A}_{\mathbf{z}'}^{(3)}$ for some $\mathbf{z}' \in \mathbb{R}[\mathcal{G}]$, if and only if there $\mathbf{z}' = \mathbf{z}^g$ for some g in \mathcal{G} .*

Unfortunately Kakarala’s proof is non-constructive, and we still do not know how to run Algorithm SM-II.1 for general groups (but see [17] for an algorithm that works for the group of all 2×2 unitary matrices with determinant +1). The proof of Theorem SM-II.1 relies on Tannaka-Krein duality (Proposition 1, [8], p. 199).

Note the following important points. Note Theorem SM-II.1 only requires condition (SM-II.13) (*i.e.*, does not require the labeling $\varrho_1, \varrho_2, \varrho_3, \dots$), whereby one seems to be able to satisfy it by slight perturbation of \mathbf{a} . This is mis-leading, as Kondor pointed out [18], pp. 89-90, for extensions as in (SM-II.5), *i.e.*, for $\mathbf{z} = \bar{\mathbf{a}}$ for $\mathbf{a} \in \mathbb{R}[\mathcal{X}]$ of general homogeneous \mathcal{G} -spaces \mathcal{X} , the condition (SM-II.13) turns out be mostly unsatisfied. While Kakarala has yet another remarkable completeness result for homogeneous spaces (see [17], Theorems 4.6 & 4.7), however as Kondor also pointed out (p. 91), this result applies only for elements in $\mathbb{R}[\mathcal{G}]$ that are constant under *right cosets* of \mathcal{S} (or invariant under left \mathcal{S} -translation as in [17]), as opposed to our definition (SM-II.5) which makes extensions constant over *left cosets* of \mathcal{S} . Hence Kakarala’s result does not apply exactly to our setup.

In conclusion, there exists some powerful results (*e.g.*, Algorithm SM-II.1 and Theorem SM-II.1) developed for the triple-correlation. However for general groups, there is room to improve these results, especially worthwhile would be a completeness result for homogeneous spaces for extensions as defined in (SM-II.5).